Spacelabs Medical

CLINICAL INFORMATION AND TECHNOLOGY SERIES

CLINICAL INFORMATION SYSTEMS

VOLUME 2: CLINICAL INFORMATION SYSTEMS OUTSIDE OF HOSPITALS
VOLUME 3: INFORMATION SYSTEMS TECHNOLOGY

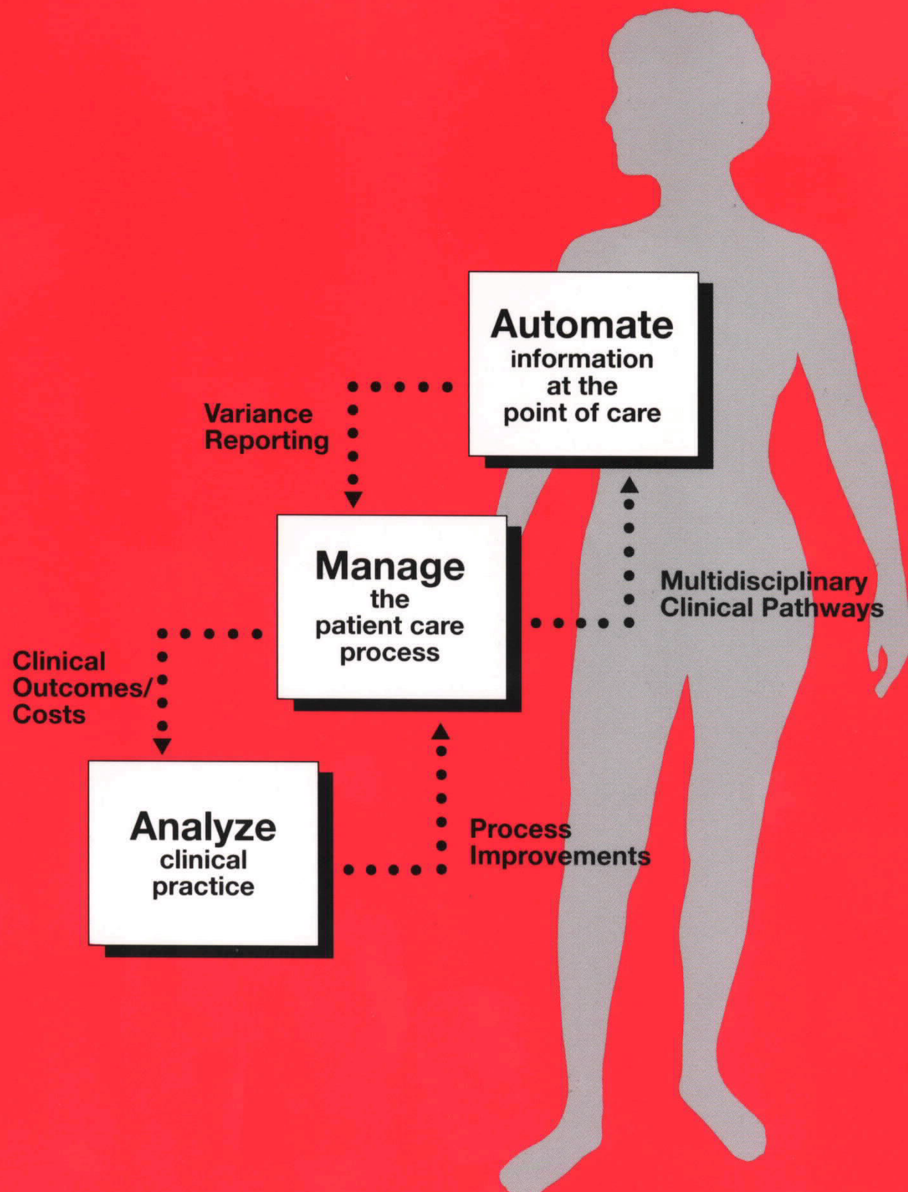

CLINICAL INFORMATION SYSTEMS

VOLUME TWO

Marion J. Ball, EdD
University of Maryland at Baltimore
100 N. Greene St.
Baltimore, MD 21201

First Consulting Group
2 Hamill Rd.
Quadrangle West 359
Baltimore, MD 21210

Michael Bourke, PhD
Houston Baptist University
7502 Fondren Rd.
Houston, TX 77074

Judith V. Douglas, MHS
First Consulting Group
2 Hamill Rd.
Quadrangle West 359
Baltimore, MD 21210

Jim Fackler, MD
Johns Hopkins Hospital
Pediatric ICU
600 North Wolfe Street
Baltimore, MD 21287

Issac Kohane, MD, PhD
Children's Hospital
300 Longwood Avenue
Boston, MA 02115

Christopher Lindberg
201 West 16th St.
Hays, KS 67601

VOLUME THREE

Marion J. Ball, EdD
University of Maryland at Baltimore
100 N. Greene St.
Baltimore, MD 21201

First Consulting Group
2 Hamill Rd.
Quadrangle West 359
Baltimore, MD 21210

Michael Bourke, PhD
Houston Baptist University
7502 Fondren Rd.
Houston, TX 77074

Judith V. Douglas, MHS
First Consulting Group
2 Hamill Rd.
Quadrangle West 359
Baltimore, MD 21210

DSA Systems, Inc.
1110 North Glebe Rd., Suite 400
Arlington, VA 22201

This book is part of the Spacelabs Medical Clinical Information & Technology Book Series for biomedical and clinical professionals. The series is an educational service of Spacelabs Medical, a leading provider of patient monitoring and clinical information systems.

Published by Spacelabs Medical, Inc., Redmond, Washington, U.S.A.

Printed in the United States

ISBN l-882588-53-3
ISBN 1-882588-54-l

TABLE OF CONTENTS

VOLUME 2

INFORMATION SYSTEMS OUTSIDE OF HOSPITALS

INTRODUCTION

by: Marion J. Ball, EdD and Judith V. Douglas, MHS

As the first volume in this series demonstrates, clinical information systems (CIS) are changing the way healthcare is delivered in hospitals, which today are also key players in the move toward integrated delivery systems. These changes reach beyond hospital walls to satellite clinics, providers' offices, and even patients' homes. Increasingly, CIS are giving clinicians access to information when, where, and how they need it.

The growth of integrated delivery systems in the managed care environment places new emphasis on care delivered by multidisciplinary teams in a variety of settings, Such teams are now common in acute care settings and are becoming more so in other settings as well. For example, managed behavioral healthcare networks extend beyond psychiatrists and psychologists to include social workers, addiction counselors, and other therapists, Nurse-managed centers offer care in poor urban neighborhoods, providing cost-effective services while assuring access to physician and hospital resources as and when needed.

The move from medical record to computerized patient record (CPR) reflects more than the move towards automation. It represents profound changes in our healthcare system, changes we are now just beginning to see. As Larry Weed, the father of the problem-oriented medical record, noted over twenty years ago, the patient is the true owner of his health record — not the physician or the hospital. Today, with managed care giving rise to integrated delivery systems, patients seek care from a wide range of providers throughout the community.

The contributors to this volume speak to these changes. Lubinski examines the value that CIS offer physicians in their practice, with a special emphasis on the managed care environment. Even today, the physician's office — where patients receive much of their ongoing care, including wellness services and disease management — is often not automated in any significant clinical aspects. With the growth of management service organizations and strategic information systems planning, this can and will change.

Lindberg offers a view into the future, with his discussion of a cutting-edge telemedicine program in the American midwest, where patients stay in their own homes and receive care through the use of video monitors and other technology. As home healthcare continues to grow, programs such as the one Lindberg describes will become more common. The evaluation approach taken by Lindberg is critical in ensuring that these programs do in fact make care more accessible and convenient.

Like telemedicine, the Internet was until recently a futuristic notion. Today it is a reality that affects many sectors and lives. Fackler and Kohane address the capabilities of the Internet and community health information networks (CHINs). The Internet has grown exponentially over the past few years, and we are beginning to understand how we can use the capabilities it provides to improve healthcare.

Lubinski, Lindberg, Fackler, and Kohane offer their views of CIS outside the hospital walls. Patients and providers alike will be empowered by newly acquired and enhanced access to health information. Many issues remain to be addressed. What measures must be taken to ensure data security and to protect patient confidentiality and privacy? How will patients and clinicians learn to use these new technologies and, more importantly, the information they provide? What will our healthcare delivery system look like and how will its components work together?

The vision we hold for clinical informatics is that of an approach that crosses disciplines and professions and focuses on the uses of technology rather than the technology itself. We believe that clini-

cal informatics and CIS will change the very process of care, thereby serving the health of the individual patient and the population as a whole.

We conclude this brief introduction by noting the early death of Chris Lindberg, who contributed his insights to this volume. We rejoice in his presence here, while mourning the tragic loss of his talents to the informatics community he served and the family he loved.

References

1. Weed LL. Medical Record. In: *Medical Education and Patient Care.* Chicago, IL: Year Book Medical Publishers, Inc.; 1969.

1.0 CLINICAL INFORMATION SYSTEMS OUTSIDE THE WALLS: PHYSICIAN OFFICES

by: David Lubinski, MA, MBA

1.1 Introduction

The practice of medicine from the physician viewpoint is the most critical element of developing a strategy for information systems. It is at the point of care that information will have the most impact on the cost and quality of care delivered. This requires that physicians prepare themselves not only for how they will use information, but more importantly for the manner in which they will *integrate* into larger delivery systems that will assume complete risk and provide services across the entire continuum of healthcare services.

This requires a significant strategic assessment of where the physician practice wants to be in the value chain of delivering care. The higher the quality and the greater the depth of patient clinical information, the more value the physician practice will have. If the physician practice decides to have an outside party provide this level of information capability, then the value is transferred to that entity. This chapter covers the required elements, from the physician viewpoint, for building clinical information system capabilities. It will then outline a practical approach to building these capabilities.

1.2 Past and Present; The Roles of the Physician Office

The evolution of information technology in healthcare has monumentally transformed the way healthcare is delivered, affecting not only the providers and payers but the patients as well. A clinical information system can be used in three ways:

1. To collect information for data tracking and analysis — establishing care plans and efficient utilization procedures.
2. To verify eligibility, payment, and coverage — reducing paperwork and ensuring covered treatment.
3. To ensure continuity and eliminate redundancy throughout the continuum of care — improving treatment.

Most often, these methods are combined on a physician system in conjunction with how the practice is integrated with local healthcare delivery systems and managed care organizations.

The following scenarios illustrate how information systems have been used in a typical physician office environment.

1.3 Scenario 1: Physician Practice Without Integration or Advanced Automation

Without the influence of managed care organizations and the use of technology in healthcare, the typical physician office visit begins with greeting patients at the reception desk and filling out a billing form. The administrator writes in the patient's name and chart number and then inserts the form into the folder that represents the patient's medical chart. In some cases, the form is imprinted with the patient's member card by running it through a manual embosser. The patient is most likely scheduled for an appointment slot in a scheduling book. Once called into the exam room, the patient would be seen by the attending nurse who completes an initial symptom / status / review and records these impressions in the paper chart.

The physician then reads the chart upon entering the exam room. This usually is his first encounter with the chart. The physical examination then takes place, diagnostic impressions made, prescriptions written, and a code or codes marked on the billing form. Once the patient leaves the exam room, he is responsible for finding a pharmacy to fill the prescription. If the patient is covered under a prepaid health plan, the copay made at the registration desk should be the only payment for this visit. However, if this visit is to an out-of-network physician, then the rules covering out-of-network care come into play. This will likely call for one of two actions. Either the patient pays for the full office cost and then reconciles with his health plan, or he pays a greater percentage of the visit than the copay. In either case, this transaction is usually handled at the front desk.

1.4 Scenario 2: Physician Practice with Advanced Automation

At the time the patient is scheduled for an appointment, an insurance eligibility determination is made electronically by entering the patient's member number. The result is a verification of eligibility, copay requirements, and deductible obligations. This information is verified at the time the patient shows up for the appointment by swiping the patient's card into a card reader, much like the readers used by an automatic teller machine (ATM) or a credit card transaction machine.

The computer system then queues the patient's medical record for review before the patient is called into the exam room. The record indicates medications prescribed, diagnoses, treatment, and outcomes. This record also contains information on the patient's prior allergic reactions or other problems and may even contain information generated when the patient was seen in the emergency room of the health system hospital. Although the patient was likely seen by a different physician, this information as well as previous outpatient information is made available to the treating physician.

Upon preparing to enter the exam room, the physician will review the electronic medical record on a flat screen outside the patient's room. The physician is presented with a color chart that indicates vital sign trends, most recent medical encounters (including the ER), and past medications prescribed. He then enters the room to greet the patient and address the medical condition responsible for this visit.

After reviewing the electronic medical record, the physician focuses on the patient's current condition, listens carefully to the patient, and asks additional questions, all the time working toward a current diagnosis. Major symptoms are entered into an exam room system which generates most likely diagnoses and recommends treatment protocols, including costs and efficacy for each. The physician then reviews this information with the patient, arrives at a diagnosis, and determines a treatment path. Medications, if required, are electronically ordered in the exam room for pick-up at the pharmacy counter in the clinic.

All information and data on the visit — cost for the office visit, waiting time before entering the exam room, waiting time in the exam room, treatment time, diagnosis and protocol used — are recorded for future aggregate and comparative reporting.

While this scenario may seem far off, practicing medicine with current and accurate information will greatly impact the efficiency and efficacy of the physician practice,

1.5 Barriers to Overcome

Despite the obvious advantages of implementing a networked clinical information system, many physicians and physician groups have fallen woefully behind in using technology and organizing to optimize their inherent strengths. While payers and large hospital providers have gained market momentum and strength by controlling the delivery of healthcare through reimbursement structures, strategic alliances, and implementation of networked technologies, physicians themselves have been one of the primary obstacles to physician group cooperation.

This is due to several key factors. Many of the changes that have occurred in the last five years have turned risk and accountability away from employers and insurers to providers and patients. This alienated many physicians who felt they were moving away from treating patients to becoming "administrative physicians". Many physicians felt that the demand for tremendous amounts of paperwork and accountability for each clinical decision and treatment plan called their practices and independence into question. At the same time, physicians were being compensated for non-patient care activities. Inevitably, physicians resented the growing amount of external control over their practice and practice methods. As in any existing situation, the concept of change-and in this case it means holistic change in practice style and management— is never easy to implement. All of this has caused an enormous amount of upheaval within the physician community.

Thus, while the demand for shared information and connectivity between providers and payers has grown, those same market and technological changes have alienated many physicians and physician groups. This factionalized segment, however, has a great need to streamline access and storage of patient and treatment information to benefit both patients and themselves.

1.6 Practice Management

Recognizing the need to reorganize practice management and healthcare delivery around clinical systems, physicians have made major strides in forming more efficient organizational structures. This is a critical element to managing care more effectively and within guidelines established by payers.

If the physician is at the core of the continuum of care, then the clinical record maintained by the practice must also be the core clinical record. The clinical record must be a dynamic document, capable of sharing information with other information systems across the continuum of care. The practice management system must support the necessary and appropriate security and confidentiality processes for sharing any patient information. This capability must be able to adapt to the various information systems that make up the network covering each patient. It is this interaction among network users that currently consumes so much time due to manual and inefficient processes.

Within an integrated delivery system (IDS), physician groups will most often interact with a designated management service organization (MSO). The MSO provides the administrative function of the physician practice, much like a physician office manager. The MSO function may exist within the IDS; it may be tied to the payer; or it may reside within an independent organization.

Because of this, not all MSOs have the same objectives or provide the same services. Typical functions include billing and collections, patient scheduling, accounting, non-physician staff management, materials purchasing, facility management, and payer/patient relations. Some MSOs provide advanced clinical care management, managed care contract management, credentialing, information integration and communications in varying combinations. Three major types of services directly impact the applications to be considered by the MSO: physician practice management systems (PPMSs), managed care systems, and contract management systems. All MSOs must meet the needs of the physician practice.

Because of the relationship between the MSO and the physician practice, it is important for the practice or group manager to analyze and assess the related MSO business plan. This assists the practice in two ways. First, the plan outlines how services will be added over time, illustrating where and how the affiliated network members intend to expand. This will directly impact the application functionality and the integration/communication technology deployed. Second, the MSO plan assists the physician practice or group in acquiring or adapting its own technology so that it is consistent and appropriate for the market as well as its own practice and patients.

Whether the physician practice will be managed through an MSO, through another organization, or remain independent, a clear business strategy must be established. This strategy will guide the development of the physician practice and drive the information systems strategy. The business strategy will address such things as the degree to which the practice will assume risk for capitated contracts, the nature of relationships with single or multiple payers, single or multiple IDSs, and other practices.

1.7 Strategic Information Systems Planning

These internal and external factors will directly impact the applications and technologies needed and will guide the group or physician practice toward a progressive and appropriate Information System (IS) plan. A strategic information systems plan (SISP) guides the deployment of applications and technology appropriate for a specific practice or networked alliance. It defines the technology architecture for systems, networks and communications, applications and operations. More importantly, an SISP defines the functional requirements in light of the business strategies of the organiza-

tion. These functional requirements are then prioritized, costs projected, and scheduled into a three- to five- year tactical implementation plan.

1.8 Defining Functional Requirements

Physicians served by an MSO are more than likely to already have a physician practice management system (PPMS). An understanding of how these systems are used, the proficiency of the office staff, and the degree to which the system delivers results need to be factored into system requirements. One of the major factors influencing functional requirements is the degree of clinical care management required by physicians. This can be heavily influenced by the sequencing of advanced clinical care. The following are various functional areas of a PPMS:

- Patient Registration. Maintains basic patient demographic information, employer information, and historical insurance information on-line; generates a patient identification number.
- Charge Posting. Accepts basic information regarding each patient visit (encounter number, patient number, care center, procedure codes, etc.) while posting charges on the patient's account.
- Patient and Insurance Billing. Supports all types of insurance billing and their requirements. Must allow capture of insurance carrier coverage information and pre-screening of services and charges.
- Accounts Receivable (A/R). Generates patient statements; performs A / R follow-up, collections, and bad debt; processes payments. A PPMS must allow for the A/R management of multiple companies and also have the capability to provide separate A/R management by location or provider.
- Referral Authorization Management/Managed Care. Retrieves authorization information and tracks it for referrals, making sure that the provider group will be reimbursed.
- Claims Management. Performs routine examinations of claims in order to determine eligibility, coverage of services, and plan liability.
- Provider Maintenance. Maintains basic demographic information for each provider in the practice. Using various indicators (e.g., utilization management, quality management, member satisfaction levels, clinical performance and administrative management), a PPMS should be able to analyze a provider's performance and report it on empirical standards and comparisons to peer groups in the provider network.
- Appointment Scheduling. Provides capabilities to create physician schedule templates and monthly master schedules. Should allow for scheduling of rooms, equipment, and other resources.
- Patient Services Management. Provides the ability to enter and track patient service issues such as complaints, suggestions, and commendations. Should support survey activities and the tracking of patient education materials.
- Chart Tracking. Provides the ability to maintain medical records at numerous operating locations. Assures third party payers that appropriate and medically necessary care was provided to the plan member and properly billed.

- Electronic Medical Record. Designed primarily to document care given to individual patients, facilitate the sharing of patient inform&ion with other care providers, and maintain a database of patient information for administrative and clinical management.
- Electronic Data Interchange (EDI). Provides **the** capability for direct system-to-system transmission of data.
- Integration Requirements. Must address four domains heavily influencing integration: scheduling, eligibility, claims, and clinical records. Needs to consider electronic communications for fundamental applications such as enterprise-wide mail and Internet access.
- Scheduling. Must allow access to multiple scheduling systems through an integrated graphical user interface (GUI) or through emulation. Although the physician's office is often the point of contact, in some IDSs the MSC handles centralized scheduling for all services. This is especially true with the IDS whose focus has transitioned to ambulatory care where the largest volume of scheduling is that of outpatients. Other significant scheduling activity is related to accessing hospital or clinic order entry systems.
- Eligibility. Verifies eligibility on-line when services are scheduled or, at the latest, immediately prior to services being delivered. This may necessitate integration with the payer's systems and likely the IDS enterprise-wide electronic network. Eligibility determination also requires communication with managed care systems that may also be provided by the MSO or located at a third party administrator.
- Claims. Consolidates and integrates claims processing. This offers many IDSs greater volume discounts as well as ease of managing payer volume. Physician claims differ markedly from hospital claims, which tend to be larger dollar amounts with copay arrangements that are not point-of-service specific.
- Clinical Records. Support the flow of selected clinical information to centralized data repositories that collect, store, and support analysis of large amounts of clinical, financial, and statistical information. As the market moves to managed care, IDSs are creating such repositories to allow them to measure outcomes, quality, and cost from the capitated perspective. Increasingly, physicians are requesting access to such repositories as a-base for clinical decisions. Given that the data can include diagnostic test results, consultant interpretations, etc., non-technical issues related to privacy, confidentiality, and security must be addressed.

1.9 Selecting Partners for Technology and Operating Services

Traditionally, healthcare providers have maintained responsibility for operating information systems with employed staff. However, many vendors now offer operations and technology services, staffing, and maintenance services on a contractual basis. Deciding which option to select is dependent upon several internal and external factors: however, one thing is clear, it is vital to select a partner who can meet both immediate needs as well as the related peripheral and future demands that inevitably arise.

Once an organization has determined that it needs to acquire an information system, the first step would be to invest time and energy in completing a carefully considered system selection process. Although the functionality of a system is a vital aspect of the selection process, there are many other factors to be considered before a partnership is made. In selecting a system, the physician practice should consider evaluating the vendor according to a number of major criteria.

1.9.1 Vendor Reputation and Company Philosophy

A vendor must possess a solid reputation within the healthcare community. Its business philosophy should be consistent with that of its customers so that their relationship is a true partnership rather than merely buyer seller. In addition, the vendor should demonstrate significant stability to assure future presence and strength in the changing healthcare and information systems market. This also includes a willingness to change over time to accommodate healthcare trends. To evaluate these factors, the vendor's history should be assessed. Information on a vendor such as whether it is a wholly owned company or a subsidiary of a larger organization and its business stability for a period of several years is critical in establishing a partnership.

1.9.2 Vendor Resources

The vendor must have sufficient resources to support the applications provided. This includes daily maintenance as well as sufficient investment in future modifications and upgrades. The percentage of revenue dollars spent on research and development is a good indicator of the vendor's commitment to change with the industry. The demographics of the total staff in the company is also a good indicator of the focus of the company. For example, if a vendor has numerous marketing staff and very few development and support staff, it may be an indication that the vendor is more focused on current sales than customer support.

In assessing future product viability, it is necessary to review past and future modifications. Who determined what modifications would occur? How long did those modifications take? How were they implemented? How effective and bug-free were the modifications? Answers to this type of question are often good indicators of the vendor's competency and effectiveness.

1.9.3 Vendor References

Reference checks with the vendor's current clients are also helpful in gaining knowledge of the vendor and its reputation. These should include vendor provided references as well as non-vendor provided references. The interviews should be conducted using a predetermined survey that covers all issues of concern. The survey instrument should be flexible enough to encourage free discussion of issues but structured enough to get comparable data from numerous sources and on competing vendors.

1.9.4 System Pricing

In addition to vendor specific issues, it is important to consider how a vendor derives its pricing. Are prices based on patient volume, total users, or other factors? Are prices standard or customizable? What charges are recurring versus one-time? How does the vendor's pricing differ from their direct competitor's? Knowing and understanding these pricing nuances will assist in negotiations as well as help to assure proper and fair product costs.

1.10 System Technology

When evaluating a vendor, the technology that has been used in developing and supporting the product is an important area of concern. An understanding of all components involved is essential, including programming language, operating system, database structure, hardware platform, and the communication devices used or required by the product. Assessing current technology trends in **the** global marketplace can help determine whether a vendor is above or below the technology curve. The physician practice should also analyze how the vendor's technological environment will affect any systems they already have in place.

1.11 System Functionality

The final area of concern when evaluating a vendor and the vendor's product is the functionality of the product being proposed. A sound approach is to create a list of expected functions, ranking them into categories: mandatory requirements, desired requirements, and optional requirements. This assists in focusing the selection process on what is necessary versus those functions that are merely appealing.

It is critical to identify the implications of each vendor's methodology and the potential impact of the **system's** limitations or robustness on the practice's business plan. Methodology assessment should include developing "scripted scenarios" specific to the organization's line of business and having the vendor demonstrate the required functionality under the scenario rather than accepting a "we can do that" attitude.

1.72 Conclusion

Physicians hold the key to effective delivery of patient care. By understanding and deploying tools that support the effective and efficient practice of medicine, physicians can improve the quality and the cost of care. Physician practices, however small, can increase the value of their service to their patients-and to the delivery networks. The advanced clinical support **systems** currently available are most definitely within the reach of physician practices.

2.0 IMPLEMENTATION OF IN-HOME **TELEMEDICINE**

by: Christopher Lindberg

2.1 *An Overview*

The history of telemedicine is described well by Perednia and Allen.’ They describe the concept of telemedicine as “The use of electronic signals to transfer information from one place to another. Telemedicine systems can be characterized by the type of information sent (such as radiographs or clinical findings) and by the means used to transmit it.”

Although telemedicine is experiencing immense popularity these days, the concept has been around since 1959 when the Nebraska Psychiatric Institute in Omaha, Nebraska, utilized a microwave link to connect with the state mental hospital 112 miles away.* Simultaneously, Jutra implemented teleradiology in Montreal, Quebec, sending telefluoroscopic examinations over a coaxial cable transmission.[3] Teleradiology is an excellent application of telemedicine utilized daily in numerous sites around the world.[4]

During the 1970s and 1980s, activity continued to a limited extent throughout the world, with projects in North America and Australia. One of the most noted was the Space Technology Applied to Rural Papago Advanced Health Care (STARPAHC) project of the National Aeronautics and Space Administration (NASA) in southern Arizona.[5] Another early project provided emergency medicine capabilities at Logan Airport in Boston, Massachusetts.[6]

In the past ten years, the most significant advancement in telemedicine has been the development of image digitization technology. Data compression technique allows transmissions of large amounts of video information with the use of relatively smaller bandwidth. An uncompressed video signal that normally transmits at 90 bits per second can now be transmitted digitally at 884,000 bits per second. The key here is the ability to transmit high resolution, full motion audio-video images over long distances. The newest advances come with the use of fiberoptics and coaxial cables that support high-speed data transmissions, and with the use of efficient data communication technology such as asynchronous transfer mode (ATM). These advances are also known for their cost efficiency.

In 1993 roughly 2,250 patients were seen through teleconsultation throughout the United States and Canada. By 1995, 60 to 70 interactive television programs were being created or implemented across the United States. Today telemedicine networks are operational in Kansas, Montana, Georgia, and Oregon. These systems are used for clinical, medical, and administrative purposes. Regrettably, to date the vast majority of telemedicine programs in the United States have low utilization rates. Actual patient interaction numbers are quite low in comparison to the cost of equipment and number of sites.

Currently 13 federal agencies are involved in the funding of telemedicine projects, including the Office of Rural Health Policy, Health Care Financing Administration, Department of Commerce, and Department of Agriculture. Approximately $100 million in funds are available through these federal sources. The fate of many

telemedicine projects often depends on the availability and continuation of federal funding. Many of the more successful state telemedicine projects were launched with federal funds and have since gained the support of other medical, educational, and administrative entities. It appears that the success of a telemedicine program is directly proportional to the continued commitment of the organization in which it resides and that organization's actual need to conquer distance.

The evaluation of telemedicine resulting in solid outcome data remains, for the most part, elusive. Clinical descriptions of effectiveness tend to be anecdotal rather than analytical. Stories relate improved patient conditions and chronicle emergency scenarios; little hard outcome data backs them up. Telemedicine program directors can know intuitively that patient health is better. However, they often find their efforts to document these and other improvements confounded by unclear research criteria, inconsistent patient selection, and low usage rates. Clearly, evaluation requires a research model for tracking the accuracy, reliability, and clinical utility of telemedicine as a primary diagnostic or therapeutic tool.

Improvements in technology and evaluation will foster the movement of telemedicine into the home. With increasing numbers of elderly patients and efforts to keep down hospital and nursing home costs, telemedicine must focus on in-home technologies and delivery. Nurses will play a crucial leading role in providing affordable, accessible in-home telemedicine used daily for clinical and therapeutic applications.

The case study that follows relates the experiences of Hays Medical Center in northwest Kansas in a demonstration project to improve both access to and quality of care. Its mission was to reach and to serve patients over long distances. It is also intended to provide a research model for evaluating telemedicine and to resolve the areas of difficulty mentioned above. Significantly, it was never designed to "show off" a new technological gadget. Its purpose was and is to improve care.

2.2 Case Study: Answering an Elderly Patient's Needs in Rural Kansas

Hays Medical Center is located in northwest Kansas and is the only secondary care center in the area serving 27 rural counties. Telemedicine is therefore viewed as a necessity rather than just another technology gadget. Currently rural hospitals in the area use teleradiology and store forward technology for teleoncology. They are linked by high speed data lines (T1 trunk lines) and interactive television (ITV), which ties rural emergency rooms to Kansas University Medical Center. The latest addition to the telemedicine program is the Interactive Home Health Program which has brought telemedicine to the elderly patients' homes.

The project was started with a grant from the Kansas Health Foundation, and most recently expanded by a major award from the NTIA program through the Department of Commerce. The grant program designates Hays, Atwood, and Lawrence, all in Kansas, as sites, with 15 patients at each site participating. In addition, a site will be included at the Veterans Administration (VA) Hospital in Kansas City, Missouri. This will allow a comparison of the use of telemedicine in rural versus urban settings.

2.2.1 Objectives

The central objective of the demonstration grant project will be to find an alternative and effective means of providing home healthcare for the elderly and disabled. It proposes to address the problems of telemedicine access to the home, premature institutionalization, and declining physical and mental conditions in the elderly patients targeted by the program.

Table 2.1 - Lists criteria used for selecting patients for the telemedicine program at two of the four rural sites.

Inclusionary Criteria
The patient is under the care of a physician and the physician will approve the plan of care. If the patient is diagnosed with a mental health condition of delusions, paranoia, psychosis or active alcoholism/drug addiction, the patient's mental health provider will approve the plan of care.
The patient resides within the specified geographic area.
The patient has two of the following:
1. Two or more chronic medical problems.
2. Four or more medications prescribed by a physician.
3. Two or more hospitalizations in the past 12 months.
4. Been referred for other supportive services in the past six months (i.e., nursing home, assisted living, day care, home care).
The patient and/or the patient's family will give informed consent for the telemedicine service.
Exclusionary Criteria
The patient has a terminal condition and is expected to live less than 6 months.
The patient requires daily skilled invasive procedures which must be completed by a nurse.
Last Level Exclusionary Criteria
Patients who will have home visits.
The patient is able to follow directions and responds to directions and/or questions properly, or if the caregiver is available during times of service.
The patient is able (either independently or with assistance) to position self in front of the monitor (monitors will be installed to best accommodate the patient).

2.2.2 Patients

The patients currently receiving care on the system represent a variety of diseases and conditions consistent with the elderly and disabled who are totally or partially homebound. Examples of these conditions include the following: chronic obstructive pulmonary disease (COPD), diabetes, Parkinson's disease, chronic depression, emphysema, etc. At the time of this writing, there were 38 patients on-line being seen weekly among three sites, and the VA hospital site is scheduled to be operational shortly with 15 additional patients. Plans call for patients to be seen by the telemedicine nurse daily or at scheduled intervals during the week. In the event the patient requires blood tests or change of a wound dressing, the telemedicine nurse would either conduct a home visit or contact the home health agency serving the patient. However, in order to properly assess the impact of telemedicine on patient outcomes, it is important not to have an outside home care service visit a telemedicine patient.

2.2.3 A Sample of Patients Selected for the Program

- 56-year-old female, diabetic, bipolar disorder, high blood pressure, manic-depressive. Has been hospitalized multiple times for either hyperglycemia (too much blood sugar) or hypoglycemia (too low blood sugar). In addition, psychiatric problems usually follow these episodes. The telemedicine nurse has been monitoring her insulin shots and glucose levels every morning, noon and night. At the start of the visits, her blood sugar range was 100 to 500. In the second week of her care plan, the range was 100 to 375. She is now averaging 225. There have been no hospitalizations; she is accustomed to the daily visits and is now awake early and consistently.
- G-year-old male patient, cancer with voice box removed, has serious agitation and aggressive behavior problems. The manager of the facility decided that the patient would have to vacate the apartment and be put into a state nursing home. The telemedicine nurse now sees the patient morning and night to be sure that the proper medications are taken. There have been no outbursts or problems for two weeks, and the patient will now be allowed to stay in his apartment and will not be moved into a nursing home.
- 72-year-old male with multiple conditions, had required a home health aide to dress and groom him. He is now up every morning dressed and groomed for his visit with the telemedicine nurse.

In general, many of the female patients who were not usually awake in the morning are now dressed, have their makeup on, hair combed and are ready for the interaction. They are beginning to treat the interaction process as an actual event that they look forward to each day.

2.2.4 Providers

Currently there are four registered nurses serving four sites. There has been significant nursing time spent at these sites on program education, equipment set-up, installation, and patient assessment. The providers in the project consist of two rural hospitals, an urban nursing home facility, and a large Veterans Administration Hospital in Kansas City, Missouri. The VA site will have its own base station, as do the rest, and a separate nurse for care. Each system is limited to its own local cable system at this point. However, the VA hospital will be using ISDN because of cable negotiation problems. As of this writing, it is too early in the evaluation process to describe provider satisfaction.

Nurses see patients at scheduled times during the week. Because of the variety of patient conditions, they are seen from one to three times daily to one to three times weekly. The standard for most patients is that they are being seen only by the telemedicine nurse and, when a home intervention is required, the nurses will conduct the visit or assign a home health nurse from the institution to make the visit. Usually the visit entails drawing blood, bandaging a wound, changing a dressing, etc. In rural sites there are limited numbers of elderly prepared to sign up for the project who meet the patient section criteria. It was for that reason that is was necessary to make the patient selection broad.

2.2.5 Equipment

The system provides full video and audio interactive capacity from Hays Medical Center to elderly patients at home . In addition, the system allows for home health outpatient medical record data to be documented and utilized. The system also allows monitoring of blood pressure, medication, diabetic condition, diet, hygiene, and mental health status. The two basic pieces of equipment for the system are the base station from which the nurse transmits and a unit in the patient's home.

The primary requisites in the design of the ResourceLink base station were to:

- Allow for acceptance of multiple transmission mediums (cable, ISDN, POTS, etc.)
- Build in flexibility to move with emerging technological changes

The base station consists of a minimum of two components:

- High speed computer with a color monitor and ZIP drive to manage the patient visits and data
- Video camera, color monitor, and headset to allow interactive patient audio/ video communication

The patient management software is a relational database built on Microsoft's Access database package to allow for upward integration with Windows NT or UNIX and application to other healthcare system databases. It includes the following components:

- A complete patient file including the Health Care Financing Administration (HCFA) assessment (information on the patient's social support, functional capacity, current medications, a review of health and nutritional status, body systems, communication level, cognitive and behavioral status, and a review of the home environment including home and financial management)
- Clinical protocols and pathways
- Physician's orders, including acceptable parameters for vital signs and red flagging if patient data exceed those limits
- Task lists, documentation standards, and visit reports
- Resource information

Comparative data analysis from the software will allow for quality improvement leading to even more effective management and improved patient and fiscal outcomes.

Since patients are already familiar and accepting of television, this medium was selected. The patient unit consists of a 13" color television with a video camera affixed to the top. The unit has no dials, buttons, or keyboards and all proprietary modifications are contained within the monitor. After receiving a "beep" from the unit two minutes before the agreed upon appointment time, the patient simply sits in front of the unit. Simultaneously, the patient's picture and audio are transmitted to the nurse's base station.

2.2.6 Obstacles

As with any innovative project, there have been obstacles — some anticipated and some surprising. Providing high quality audio and visual transmissions proved to be more difficult technically than expected. Project personnel had originally been concerned about providing visual images with sufficient clarity to be readable to the telemedicine nurse at the central site. Doing so was no small task given the need to read syringes, medication labels, glucometers and blood pressure devices. However, the challenge of providing good quality audio transmission turned out to be more problematic. Initially, there was either no sound at all or very poor quality sound. Compounding the problem was the fact that many of the patients already had hearing problems. Project personnel eventually determined that the audio loop in the software was causing feedback. The problem with the software was corrected, vastly improving the audio quality.

Kansas University Medical Center has completed the evaluation research model that is entered into a computer and used at all four sites. Although the questions in the model are relevant and well thought out, the manner in which the nurse is expected to answer them is too time-consuming. It appears to be a software problem rather than any difficulty with the questions themselves.

The telemedicine nurse at Hays has experienced some alienation and isolation due to the nature of the job. It has been very difficult to explain the process for success and anticipate problems with such a completely new idea and project. The nurses continue to work hard to learn new ways to utilize the software and make the interactions meaningful for the patients. Central problems occurred in following nursing procedures and protocols, An example was interpreting care plans, physician orders and actual patient care responsibility versus technical problems, or not being sure

which procedures to follow when dealing with problem patients. Much of this stems from standard home health policy and procedure versus utilization of the new system.

The interactive telemedicine project creates the phenomenon of the nurse having to be flexible, create care plans, and apply changes for some normal procedures of nursing. We solved many problems by scheduling the telemedicine nurse to see the VP of Nursing on a regular basis. This allows the telemedicine nurse to discuss specific nursing problems, change focus on the home unit, change lighting or solve a patient problem that would require a nurse at the base station and at the home unit itself.

2.2.7 Reimbursement

The Medical Center's home health department makes 1,025 home health visits per month on the average. Ellis county's population currently stands at 26,004 with 17.5% or 4,547 over the age of 65. The surrounding counties in northwest Kansas have anywhere from 9% over the age of 65 to 27% in Smith county, which has the highest rate of seniors per capita in the United States. The average home health visit runs over $60 per skilled home health nursing visit to the patient. It is estimated the in-home video service will be approximately $30 per visit. Kansas Social and Rehabilitation Services and HCFA agreed favorably for Medicare reimbursement utilizing the current system. Negotiations with the HCFA for Medicaid are still in progress. The proposed reimbursement is defined through routine visits.

2.2.8 Evaluation

The evaluation for the project is currently performed by Allen and Whitten at the Kansas University Medical Center/Information Technology Services and Research. The evaluation is still in its early stages. Conclusions and implications cannot yet be drawn, and it would be inappropriate to comment on patient outcomes and acceptance during the implementation phase due to the relatively small number of patients. The basics of the evaluation plan will cover parts 1 and 2, as described in Table 2.2.

Table 2.2 – Part 1: Descriptive Analysis.

Part 1 - Descriptive Analysis	
RQ1-	Demographics of patient base, utilization, diagnostic categories and unforeseen events and effects.
RQ2-	Intensity of nursing care required by patient base.
RQ3-	Describe mental, physical, and emotional status of participants.
RQ4-	Chart types, frequency, dosages, and costs of participants' medications.
RQ5-	Chart services provided (types and modalities of nursing services provided).

Table 2.3 - Part 2: Satisfaction.

Part 2 - Satisfaction	
RQ1-	Measure levels of patient's satisfaction in receiving home telehealth services.
At each interaction, the following research data are collected:	
1.	Vital signs: BP, pulse, temperature, blood glucose. Any measurement outside normal parameters?
2.	Medication Administration: Were all medications taken as ordered without adverse effect?
3.	Mental Status: Alertness, affect, orientation. Has there been a change in mental status?
4.	Functional Status: Mobility, activity tolerance, activities of daily living (ADL) and / or instrumental activities of daily living (IADL). Has there been a change in functional status?
5.	Nutritional Status: Is nutritional intake adequate and compliant with restrictions?
6.	Overall Health Status: Has there been a significant change in overall health status?
7.	Weight Monitoring: Has the patient's weight changed beyond normal parameters?
8.	Patient Teaching: Did the patient receive instruction related to diagnosis and / or treatment?
9.	Change: Was there a medication change since the last visit?
10.	Change: Was there a diagnosis change since the last visit?
11.	Utilization Data: Emergency room visit? Yes or No.
12.	Utilization Data: Doctor / other outpatient visit? Yes or No.
13.	Utilization Data: Admission to hospital? Yes or No. Length of stay?
14.	Utilization Data: Admission to nursing home? Yes or No. Length of stay?
15.	Utilization Data: Other healthcare services received? Yes or No. Describe.
16.	Technical: Were there technical problems? Yes or No. Describe.
17.	Visit Type: Phone / on-site / televisit.
18.	Visit Type: Scheduled / unscheduled.
19.	Visit Type: Acute / chronic / both.
20.	Visit Data: Nurse / patient initiated.

2.3 Conclusions

As of this writing, the project is in its sixth month with 38 patients online at three sites; a fourth site is being installed. Evaluation data are limited, making speculation on those data premature. Project personnel have experienced many successful patient "profiles" in which the patients showed marked improvements in specific situations.

The utilization of the local cable system has been extremely effective in creating a high definition quality telemedicine audio and visual system. Project participants are encouraged by the daily use of telemedicine in the home.

2.4 Acknowledgments

The author gratefully acknowledges the insights provided by Douglas Perednia, MD, Biomedical Information Communication Center and Department of Dermatology, Oregon Health Sciences University, Portland, Oregon; and Ace Allen, MD, Cancer Center, University of Kansas Medical Center, Kansas City, Kansas. Their generosity in sharing their experiences and expertise is profoundly appreciated.

2.5 References

1. Perednia D & Allen A. Telemedicine **technology** and clinical applications. *JAMA*. 1995; 273(6):483-487.
2. Wittson CL, Affleck DC, Johnson V. The use of two-way television in group therapy. Omaha NE: Nebraska Psychiatric Institute; 1961.
3. Jutra A. Teleroentgen diagnosis by means of videotape recording. *Am J Roentgenology*. 1959;82:1099-1102.
4. Allen A, Allen D. Teleradiology 1994. *Telemedicine Today*. 1994;2(3):14-23.
5. Rayman BR. Telemedicine: military applications. *Aviat Space Enviorn Med*. 1992;63:135-137.
6. Bird KY. Cardiopulmonary frontiers: Quality health care via interactive television. *Chest*. 1972;61:204-205.

3.0 THE FUTURE: INTERNET, **TELEMEDICINE** AND CHINS

by: Jim Fackler, MD and Isaac Kohane, MD, PhD

Forty years ago digital computers were in use for large-scale data problems.* Thirty years ago, punch cards and magnetic tape were implemented to improve military healthcare delivery at lower costs.[2] Yet, the admonition "...that computers be recognized...and that...everyone learn that we are limited only by our imaginations"[3] has not been realized. So, to parrot words from many years ago and to now suggest the future of medical computing will "...eliminate communication bottlenecks...assure complete, concise, legible patient records..and...reinforce the physician's memory..."[4] seems a bit like predicting rain in the desert, eventually it will happen.

Certainly computers are ubiquitous in medicine, particularly in critical care applications. There are a variety of computer systems and some are quite sophisticated. However, implementation of even medical administrative functionality lags far behind other industries. Inventory control in hospitals is far behind grocery stores. Point-of-contact data acquisition is far behind the automobile rental industry. Process control is far behind manufacturing. Literally billions of dollars are being spent to bring medicine to the computing level which these industries consider routine. Although a medical computer revolution may seem upon us, the true revolution has yet to be realized.

There are, however, pressures which make the grandiose predictions of previous decades worth revisiting. As it was cost containment that drove grocery stores to automate inventory control, so too will cost containment drive medical automation, As the modern financial industry exists only with automation, so too will medical networks exist only with automation. Patient care will improve as a result of automation similar to the way grocery stores benefited in the form of accurate stock management and banks from fund transfers between institutions.

This chapter will examine three important phenomena in the future of clinical computing:

1. Foremost, the authors predict the Internet will revolutionize medical computing. Many substantial problems which have plagued the medical informatics community are solved within the Internet. Of paramount importance are the Internet's universal information naming scheme (i.e., hypertext transfer protocol or HTTP), extraordinarily inexpensive clients (i.e., web browsers), computing platform independence, and security.
2. Telemedicine, defined simply as the practice of medicine at a site remote from the practitioner via audio, video, and/or data linkages, will become more widespread. Again, driven by market-building networking pressures (and formally recognized and reimbursed in some states), telemedicine will see increasing utilization locally (to link office-based physicians with their hospitalized patients), regionally (to link large medical centers with rural patients), and internationally (to deliver expert advice regardless of location).
3. Less important is a discussion of community health information networks (CHINS). Many CHIN discussions center around the benefits to patient care. Yet, more realistically, the structure of these networks will be driven by market pressures.

3.1 The Internet

A substantial deterrent to the ubiquitous introduction of electronic health record systems (EHRSs) has been the belief that medical data are different. Certainly the encoding of medical text (particularly history details) is elusive, yet procedures and diagnoses have been codified for many years. If, however, one accepts the paradigm shift from viewing medical information sciences as being wholly unique, then searches for an EHRS solution can begin within other commercial data processing systems designed for other information-intensive industries. Substantial savings can be anticipated with a transition away from custom designed systems. More important, use of standard data processing systems will promote medical industry-wide communications.

There is no standard data processing system more open than the world wide web (WWW). The world wide web, or simply "the web", was conceived nearly 30 years ago to solve the problem of slow distribution of scientific knowledge. From what was once a network of no more than a handful of computers, the web has now grown in size to over ten million computers with the number of users doubling every five to eight months. Web connections which were originally limited to academic institutions are now open to worldwide access from any personal computer with a modem connection to an Internet service provider. The web is a computational platform and is independent of the underlying operating system. Developers of web browsers typi-

cally give out the software free or at a nominal cost to users, thereby making it attractive to everyone, including computer neophytes, to surf the net. Finally, new web-based technologies will prompt the development of applications that will further enhance the power of the web. For example, solutions that address Internet security and confidentiality issues will promote Internet bank transactions, stocks trades, and medical record transfers. Another useful web technology for medical applications is the Java programming language. Java allows multiple continuous processes to remain active between two computers, allowing real-time display of data such as physiologic waveforms.

Technically, the web is a collection of standard communication protocols. Hypertext transfer protocol (HTTP) is a universal naming scheme for all information on the Internet. Hypertext markup language (HTML) is a first generation markup language that allows multimedia documents to be displayed, again, on a platform independent client. Equally crucial are the HTML support links within one document to other documents. Enhancements now allow inclusion of formatted text (e.g., tables, still images, digital video, and sound, including telephone). An excellent review of the web technologies was recently published by Lowe et al.[5]

However, if the use of the Internet for EMRS were to be restricted to HTML and HTTP, the medical world would have but one more graphical user interface. To reiterate, the interface would be ubiquitous, cheap, and platform independent. However, the problem of multisite data integration would not be solved.

3.2 The Common Medical Record

Figure 3.1 illustrates the current state of disorganization. The legacy systems could represent data management systems within a single institution (e.g., an emergency department system, an operating room system and a critical care system) or could represent three hospital systems within a network. Integration of the legacy systems can be achieved by completely re-implementing the legacy systems in a common structure. There is substantial reluctance to this mechanism of data integration because of the associated costs and disruptions to functioning systems. Additional reasons for the reluctance to migrate to a single legacy system include a lack of a system suitable for the needs of all users of medical data and a lack of consensus on a common medical data architecture. As shown in Figure 3.1, translators need to be written between each system and its interface, to permit each set of users to continue to maintain their legacy systems and user interfaces. Unwieldy for integration of these systems, as each intended integration is added, the number of necessary translators required approaches x^2 (Figure 3.2).

Figure 3.1 – Integration without a common medical record (CMR).

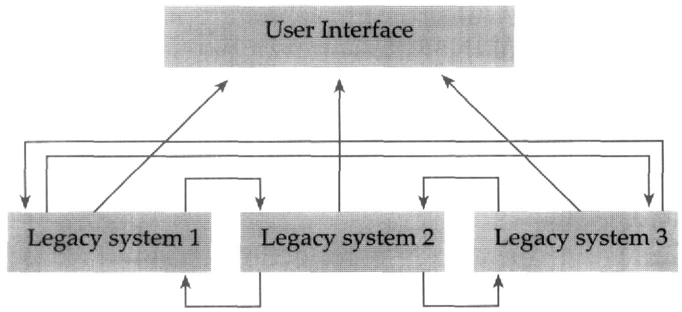

Figure 3.2 – Translators necessary to integrate "X" number of legacy systems (note logarithmic scale).

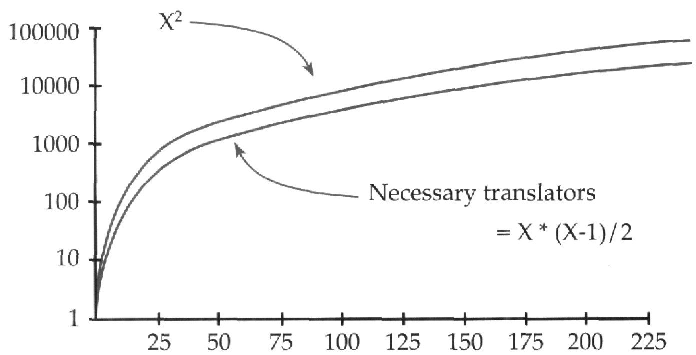

Figure 3.3 – "X" translators are necessary for integration.

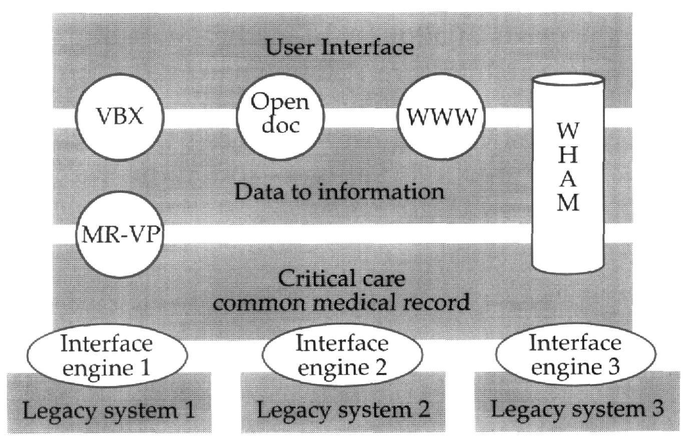

Figure 3.3 illustrates the utility of a common medical record (CMR). As a digression, the wisdom of a CMR has been obvious for decades; a myriad of private and government organizations have labored for literally thousands of hours to create the CMR. These authors suspect that the reasons why efforts have met limited success is the insufficient value added by the efforts to justify complete migration or even translation to any iteration of a CMR. However, with the reduction of the cost-to-power ratio of computers, coupled with the market pressures for data integration, the CMR as proposed in Figure 3.3 becomes viable. In contrast to the efforts at data integration without a common record model shown (Figure 3.1), integration of additional legacy systems following adoption of a CMR requires only a linear increase in interfaces. That is, the number of interfaces equals the number of legacy systems rather than the square of the systems.

Of equal importance as the interfaces between the legacy systems and the CMR, the imposition of the CMR dramatically simplifies the implementation and customization of the user interface. The abstraction layer above the common medical record labeled data to information encompasses a number of computational processes which take the raw data streams from the legacy interfaces and manipulate the data into a structure more conducive to the user interface designers. For example, while within the structure of the CMR, a chest radiograph may be embedded in radiology tables, within the user interface. A chest radiograph is more appropriately displayed within a respiratory view. Again, the abstraction between the CMR and the user interface allows for creation of standard data views (i.e., data becoming information) from which further customization is far simpler.

As shown in Figure 3.3, the web is one of the technologies which may be used to deliver information to the user. For performance purposes, it may be desirable in some circumstances to consolidate the processes into a single program such as WYSIWYG HTML authoring for medicine (WHAM) which negotiates data to information transformations directly from the CMR to the user interface.[6]

3.3 The Internet and the Common Medical Record

Figure 3.4 – Multisite data query.

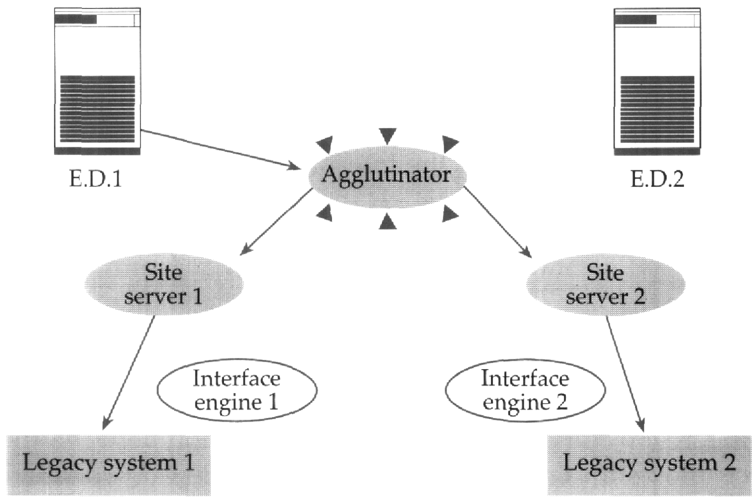

Figures 3.4 and 3.5 sketch the architecture used by the Boston EMR Collaborative for multisite aggregation of emergency department based patient data.[7] As Figure 3.3 shows, the web is used for the user interface. Once a request for information on a specific patient comes from the emergency department, it is transmitted via HTTP to the agglutinator. The agglutinator encompasses a variety of functions but in essence rebroadcasts the data request to all site servers where relevant patient data may be stored. A standard messaging language, HL7, is used for each data request and delivery. Each site which receives the HL7 request for information translates the request into the local data manipulation language appropriate for the local legacy system. The responses are translated back into HL7 messages which contain the necessary CMR information and transmitted back to the agglutinator (Figure 3.5). The agglutinator converts the HL7 to HTML for transmission to the HTTP server which forwards the information to the requesting user interface. Crucial in this architecture is the incorporation of widely accepted messaging standards (e.g., TCP/IP, HTTP, HTML, HL7, and CMR).

Figure 3.5 – Multisite data retrieval.

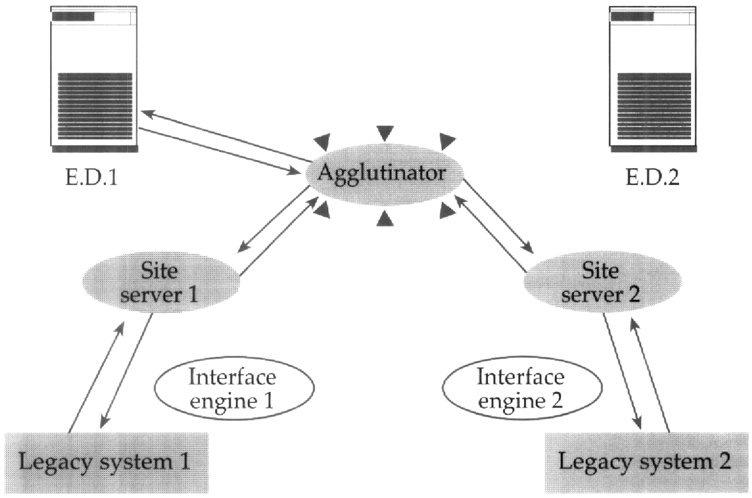

Table 3.1 – Confidentiality policy overview.

Confidential policy overview
1. The patient must maintain autonomy and control of their medical information.
2. Access to information should be role specific.
3. For transfer of any medical information:
• The patient must consent when alert.
• The patient may prospectively deny consent if they become incapacitated.
• An access/delivery log must be kept and
• The patient must be notified of all information transfers.

3.4 Internet and Security

The confidentiality of medical information which a patient entrusts to healthcare professionals is an unquestionable tenet of the traditional patient-physician relationship. That patient relationships are evolving into a patient-to-healthcare system relationship only magnifies the need for scrupulous confidentiality policies and data security standards. Further, because the Internet is correctly characterized as anarchical, the fusion of confidential medical information and the Internet immediately raises near hysterical concerns. Table 3.1 lists an overview of fundamental issues that specific confidentiality policies must address. Further, data security standards must *at least* meet current standards (Table 3.2).

Table 3.2 – Security technical solutions.

Security Technical Solutions
1. Secure communications:
• All communications with HTTP.
• Agglutinators behind firewalls.
• At least one agglutinator site.
2. Authentication:
• Source site authentication.
• Destination site authentication.
• User authentication with hardware token.
3. Patient identification by key identifiers.

However, use of the Internet will *improve* security with *existing* protocols for secure authentication and end-to-end encryption (e.g., the secure sockets layer protocol[8] and s-HTTP[9]). End-to-end encryption offers the advantage over all current systems; no matter how insecure intermediate computers or network components are, the security of data transmissions can be guaranteed secure against the most determined and costly attacks — attacks far more difficult and costly than the current non-encrypted systems. Although security lapses have been uncovered with the widespread use of web browsers in many industries, more problems are yet to be discovered. Readers interested in a comprehensive bibliography on these issues are referred to a list compiled by the National Library of Medicine (NLM) located on the web at http://gopher.nlm.nih.gov:70/00/bibs/cbm/confiden.txt.

3.5 Telemedicine

According to the National Library of Medicine, "Telemedicine is the use of telecommunications for medical diagnosis and patient care. It involves the use of telecommunications technology as a medium for the provision of medical services to sites that are at a distance from the provider. The concept encompasses everything from the use of standard telephone service through high speed, wide bandwidth transmission of digitized signals in conjunction with computers, fiber optics, satellites, and other sophisticated peripheral equipment and software"(on the Internet see http://www.nlm.nih.gov/pubs/factsheets/telemedicine.html). The NLM identified 1,643 citations in the medical literature from 1966 to 1995 on the topic of telemedicine.

A second definition of telemedicine is worth sifting through because it arises from the joint working group on telemedicine of the Department of Health and Human Services. The working group distinguishes telehealth as general healthcare services from telemedicine as individual patient care. According to the working group, clinical telemedicine is the delivery and provision of healthcare and consultative services to individual patients and the transmission of information related to care, over distance, using telecommunications technologies and incorporating the following activities:

1. Direct clinical, preventive, diagnostic, and therapeutic services and treatment, including procedures where a provider may be present with the patient, and clinical training and consultative clinical grand rounds, if used for decision making regarding the clinical care of a specific patient.
2. Consultative and follow-up services.
3. Remote monitoring, including the remote reading and interpretation of results of patient's procedures.
4. Rehabilitative services.
5. Patient education provided in the context of delivering healthcare to individuals.

Telemedicine recently received international media attention. At the 1996 Olympics, athletes were provided access to full-motion video conferencing to communicate with their friends and families. Video conferencing was also utilized to transfer radiographic and ultrasound images from the Olympic clinic and several top hospitals for rapid diagnosis.

However, if this emerging technology is allowed a classic archetype, telemedicine is most often envisioned as a remote (usually rural) video and digital connection with a central (usually academic) medical center.[10] In this scenario, telemedicine is a promising vehicle to bring healthcare to sites previously inaccessible or accessible only with substantial difficulty. A telemedicine program in Colorado and Kansas began operation in August 1995 and covers 100,000 patients spread over 20,000 square miles. This program is estimated to save 80,000 miles of travel per year.[11] In addition to the travel expenses saved, the individual saves lodging and meal expenses. Further, wages are not lost during the time typically spent traveling. If telemedicine could be used to facilitate early detection of potentially serious medical conditions, the initial investment made per year per site could result in substantial savings for the healthcare network. For example, if skin ulcers were detected in a diabetic patient, simple, local, and inexpensive therapies could prevent hospitalization and its associated expenses. The concept of telehealth is, therefore, quite relevant (see Table 3.3).

Table 3.3 – Patient charge (per day).

Facility	1992	1995*
Hospital	$1459	$1810
Skilled nursing facility	$264	$293
Home care	$75	$86

**estimated*

The most advanced telemedicine application is image processing with teleradiology being the most common current imaging application. Even in rural hospitals, about 66% of telemedicine encounters are teleradiology.[12] Microscopy images are easily transmitted via telemedicine; pathologists (and the patients waiting for the results) are thus able to avoid the lengthy delays inherent in a mail-based system. In both specialties, patient care is substantially enhanced by the ease of access to expert (remote) consultation.

The catchment of a telemedicine network is limited only by the basic communication infrastructure. Outreach programs for developing countries have recently been announced by major U.S. hospitals. The Massachusetts General Hospital, for example, is connected with two institutions within the Dominican Republic. As mentioned above, the first focus of these efforts will be to deliver support in radiology and pathology.

There is an enormous spectrum of technologies that have been labeled as telemedicine. Some telemedicine systems are as simple as electronic mail and specialty-focused bulletin boards.‘* More elaborate are systems for remote cardiology consultations with real-time bi-directional audio and video connections for transmission of physical signs and heart sounds. Dermatology consultations can be facilitated with video transmission of low power microscopy of skin lesions. Full surgical procedures have been broadcast with simultaneous, continuous, multichannel neurological monitoring signals being delivered to a third site for real-time interpretation and feedback.

3.6 Telemedicine and Critical Care

Applications of telemedicine to the critical care domain are rarely discussed, yet the roots of a telemedicine system are ubiquitous in the routine use of phone and radio consultation between emergency medical technicians and emergency rooms and/or critical care units. Data are often passed during those consultations (typically electrocardiograms). Wireless technology for simultaneous video, audio, and data exchange is available, but its cost remains prohibitive.

A possible additional application of telemedicine in critical care is the linking of skilled nursing facilities and rehabilitation facilities with a central critical care unit. Patient transfers are common between such facilities. Transfers to critical care units could, in part, be reduced if real-time video and data links were established. For facilities within a network, total costs to the network would potentially decrease.

Research and development is also active in the area of home-based telemedicine.[13] In the critical care domain, a benefit of telemedicine technologies is easily postulated in the home-ventilated patient. Linking of the electrocardiogram and respiratory monitor with continuous or intermittent pulse oximetry, coupled with the ability to see the respiratory pattern, could save emergency room visits, hospitalizations and — potentially — lives.

Home monitoring (with central supervision and consultation) of usually less complicated problems (e.g., asthma or diabetes) could also have enormous impact given the prevalence of both conditions if hospitalizations were prevented (see Table 3.3).

3.6.1 Telemedicine Impediments

Funding for telemedicine network development is coming from multiple government agencies (e.g., National Institutes of Standards and Technology, National Institutes of Health, Food and Drug Administration) as well as commercial sources. In many circumstances, startup costs are no longer an impediment.

Reimbursement for services rendered is a crucial factor in the widespread implementation of telemedicine. Legislation passed in a few states and pending in others

requires insurers to pay for services appropriately provided by telemedicine. Such a law, therefore, mandates that providers be reimbursed even without face-to-face contact between provider and patient. Telemedicine has received sufficient attention from the Health Care Financing Administration that demonstration projects are funded to reimburse consulting physicians for their time regardless of whether a Medicare patient is seen face-to-face or over a telemedicine network. Complete acceptance of this is yet to be seen.

By its nature, telemedicine is location-independent. Since telemedicine crosses geographical boundaries, state licensure has become an issue that is currently being addressed by legislation. Proponents of telemedicine have suggested a system of licensure reciprocity between states as one way to address this issue. Similarly, the geographical independence of telemedicine-delivered services presents medical liability issues which must also be addressed.

3.7 Community Health Information Networks

The belief that easy access to patient data, regardless of its location, would augment patient care cannot be argued; what could be argued, however, is the extent to which it improves care. Also arguable is the acceptable risk (e.g., confidentiality breaches) in a setting of computer-linked multi-institutional patient data. In their best light, community health information networks (CHINS) offer a mechanism for improved patient care by supporting widely accessible patient data. To reiterate from the introduction, CHINS should finally realize the admonition that computers will "...eliminate communication bottlenecks...assure complete, concise, legible patient records...and... reinforce the physician's memory..."[4]

These altruistic considerations will not, however, be the driving force toward community-wide data integration. Market pressures on network formation are enormous. Third party payers demand increasing quantities of information before reimbursement for medical services is considered. Such pressures push for CHIN development. A pressure against CHIN development is the concern that market pressure and reimbursement considerations will predominate over individual patient privacy considerations.

Possibly the first CHIN was described about 30 years ago.[2] Every inpatient and outpatient visit of military personnel was recorded on a punch card and sent to a central office for data to be stored on magnetic tape and then sent over for analysis. The information was stored for a year on 24 rolls of tape, about the size of a modern commercial movie. With the information, patients were effectively followed regardless of their location. The situation is completely analogous to current health maintenance organizations (HMOs) needing to track their covered population to effectively allocate both physical and human resources.

The above example, as well as the most commonly articulated CHIN architecture, requires delivery of data to a central system. As such, most efforts have failed. Data ownership, accuracy, security and confidentiality are each an enormous problem. It is the author's opinion that a centralized CHIN cannot work in the presence of a combination of these problems. Rather, the authors believe that CHINS are best developed by linking heterogeneous legacy systems'" (see CMR section above).

The Internet and common mal record discussions have significant implications for community health information networks. Adoption of a CMR is a solid step to-

ward linkage of disparate data systems. Further, a distributed architecture allows individual institutions to maintain control of and, therefore responsibility for, individual data sets. The reasoned fear of a single, massive, centralized database becomes less (but not incompletely) irrelevant. Finally, with an appropriately fine-grained data structure, individual patients can dissociate and independently grant access to subparts of their medical information. For example, a patient may wish to grant a CHIN access to their medication and allergy lists, but not allow access to a visit history.

3.8 Conclusion

The confluence of affordable hardware, more attention to medical software, market pressure toward network formation, necessary medical cost reductions, and the slow acceptance of automation in medicine will bring changes to healthcare computing. The authors hesitate to describe these coming changes as dramatic because that would be counter to the snail's pace of automation incorporated to date. There is little doubt, however, that the rate of change will increase. Medicine may soon reach the level of sophistication of the grocery store.

3.9 References

1. Lusted LB. Medical electronics. *N Engl J Med.* 1955;252:580-85.
2. Luykx HMC. Modernizing the management of the Air Force Medical Service. *JAMA.* 1963;622-625.
3. Caller BA. The value of computers to medicine. *JAMA.* 1960;174:2161-62.
4. Best WR. The potential role of computers in medical practice. *JAMA.* 1962;182:994-1000.
5. Lowe HJ, Lomax EC, Poonkey SE. A review of emerging Internet-based technology for the distribution of biomedical information. *J Am Med Inform Assoc.* 1996;3:1-14.
6. Hinds A, Greenspun P, Wham IK. A forms constructor for medical record access via the World Wide Web. Proceedings - *the Annual Symposium on Computer Applications in Medical Care.* 1995:116-120.
7. Van Wingerde FJ, Schindler J, Kilbridge P, et al. Using HL7 and the World Wide Web for unifying patient data from remote data databases. Proceedings - the *Annual Symposium on Computer Applications in Medical Care.* 1996:643-647.
8. Hickman KEB, Elgamal T. The SSL Protocol. Internet Draft: Netscape Communications Corporation; 1995.
9. Rescorla E, Schiffman A. The Secure Hypertext Transfer Protocol. Internet Draft: Enterprise Integration Technologies, 1994.
10. Sanders JH, Tedesco FJ. Telemedicine: bringing medical care to isolated communities. *J Med Assoc Ga.* 1993;82:237-41.
11. Zeiger D. Healthcare network project in Colorado poised for expansion. *The Denver* Post. Denver, 1996.
12. Phillips DF. Physicians put promise of telemedicine to the test: reports from rural practitioners, anesthesiologists. *JAMA.* 1996;276:267-8.

13. Engstrom I'. Telemedicine targets a rapid-growth market: Home sweet home care. *Medicine on the* **Net.** 1996;2:1-4.

14. Kohane IS, van Wingerde FJ, Fackler JC, et al. Sharing electronic medical records across multiple heterogeneous and competing institutions. Proceedings - the *Annual Symposium on Computer Applications in Medical Care.* 1996:608-612.

VOLUME 3

INFORMATIONS SYSTEMS TECHNOLOGY

INTRODUCTION

by: Marion J. Ball, EdD and Judith V. Douglas, MHS

As the first two volumes in this series document, clinical information systems (CIS) are changing the way that healthcare is delivered, in and outside the hospital. With capabilities ranging from advanced instrumentation to high-level decision support, CIS offer clinicians information when, where, and how they need it. Increasingly, CIS applications function as the mechanisms for delivering patient-centered care and for supporting the move toward the computer-based patient record (CPR).

These advances — indeed, the entire clinical informatics concept — are made possible by the technology that underlies them. The second half of the 20th century saw first the development of vacuum tube-based mainframe computers and then Pentium chip-based personal computers. We have moved from the hardware era into the software era, and stand ready to enter a new age where peopleware is paramount. In the 21st century, clinical informatics will build upon the infrastructure provided by hardware and software. High performance computing and communications are coming together to create networks capable of supporting a seemingly endless array of applications. The challenge now is to learn to use the technologies already available to us to deliver, improve, and ultimately transform healthcare.

As early as the 1970s, a small group of pioneers explored the frontiers of medical informatics, passionately believing that information technology could assist clinicians and patients alike. Their work is related in *A History of Medical Informatics* (American Medical Informatics Association, 1996), authored by one of their numbers, Morris Collen. In 1979, an international working group formulated what is in essence a definition of clinical informatics, when they concluded that "...the key is not so much the technological capability, but technological performance conjoined with medical, nursing, and administrative staff's perceived need to improve the effectiveness of their components of Healthcare."' Visionaries like Larry Weed advocated using technology to extend the human mind and help individuals link clinical problem solving to the ever-growing body of biomedical knowledge.* Today these pioneers are seeing their visions of almost three decades materializing.

Tomorrow holds still greater promise. We will see improvements in and enhancements to existing technologies, making them less intrusive and more intuitive. We also expect to see voice recognition mature and handheld devices become as ubiquitous as cellular phones.

As we enter this new age, we look to those who understand the technology to give us guidance, and to help us understand what we must do so that we can take full advantage of the power it confers. As knowledge navigators, we need to grasp the concept of how the computing engine runs and what the rules of the road are. It is this understanding that the contributors to this volume provide.

DSA Systems, Inc. explains operating systems and database technology, and Michael Bourke addresses networks, workstations and computers, and input/ output devices. Together DSA Systems and Bourke offer a solid and essential base for understanding clinical information systems. This volume complements and supports the two that precede it in this series, and merits a place on the clinical informatician's bookshelf. To DSA Systems and Bourke, our thanks.

References

1. Shannon RH (ed). *Hospital Information Systems.* North-Holland, 1979.
2. Weed LL. Medical Record. In: *Medical Education and Patient Care.* Chicago, IL: Year Book Medical Publishers, Inc.; 1969.

1.0 OPERATING SYSTEMS

by DSA Systems, Inc.

An operating system is a program that controls the execution of application programs and acts as an interface between the end user and the computer hardware. The operating system consists of several component routines. The *supervisor* manages the operating system activities. It remains in primary storage at all times and therefore is said to be main memory resident. Routines that are maintained in secondary storage and brought into primary storage when needed are known as transient routines.

The operating system performs five basic functions:

1. It schedules jobs, which involves determining the sequence in which jobs are executed.
2. It manages the hardware and software resources by loading the user's application into primary storage, which causes it to be executed, and instructing the hardware to perform the tasks required by the application.
3. It is responsible for maintaining the system security by requiring the user to enter a password.
4. It handles multi-programming which allows the scheduling and execution of multiple users' applications at one time, as well as interrupts which allow the operating system to temporarily suspend one process to allow another process to be executed.
5. It records the time usage of each system unit, such as the CPU, input and output devices, and secondary storage for record-keeping purposes for each user.

1.1 Resource Utilization

An operating system allows the user to use the resources of a computer. These resources include physical memory, disk drives, CD-ROM drives, printers and any other peripheral attached to the computer. Modern operating systems also provide the ability to share memory among applications, use the disk drive as a surrogate for physical memory (virtual memory), communicate between machines, and run more than one program at a time.

Figure 1.1 – Role of operating systems in computers.

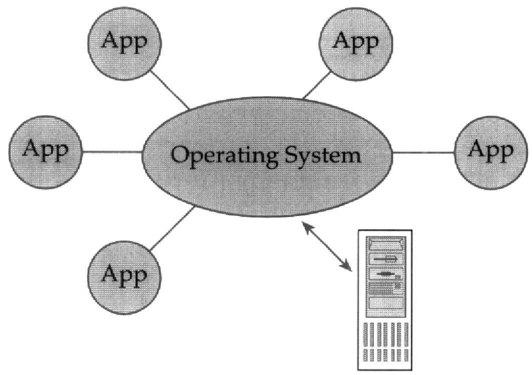

1.2 Ease of Use

An operating system allows the computer system to be more easily used by controlling the operation of all other programs for the user.

The computer system can be viewed in a hierarchical manner with the operating system making up a portion of the hierarchy. The end user interacts with applications such as Microsoft Word or Lotus 1-2-3. These applications interact with the operating system which in turn interacts with the computer hardware. This layering protects the end user from having to understand the underlying architecture of the computer.

Operating systems have evolved over time allowing for further separation of the end user from the hardware. The user interacts with the graphical user interface which sends the requested action to the operating system for execution. This protects the user from having to understand the "language" of the operating system.

The operating systems to be discussed share several characteristics:

- Each takes advantage of the full 32-bit features of the processor. Under a 32-bit operating system, programs can manipulate 32-bit chunks of data. This allows for more efficient and speedy data transfer. It also allows for more memory to be used by an application by using flat memory model addressing.
- Each takes advantage of multitasking and multithreading. ***Multitasking*** allows the operating system to execute several programs simultaneously. This permits the user to easily switch between two or more applications running on the computer. ***Multithreading*** allows an application to divide up tasks into smaller executable sections. This permits the application to do several things at once, such as loading a file while updating the user interface.

1.3 Windows 95 Operating System

As of this writing, Windows 95 is the latest version of the popular Microsoft Windows operating system. An enormous amount of software already exists for Windows 3.1 and MS-DOS. A significant effort, therefore, went into maintaining backward compatibility, allowing the existing software products to function in Win-

dows 95. In addition, Windows 95 provides enhanced features, streamlines the user interface, and improves the performance of the Windows environment.

1.3.1 User Interface

Windows 95 is a radically new, more productive working environment compared to its predecessor, Windows 3.1. The user interface has been redesigned with object-oriented characteristics. The new user-friendly interface includes context menus, a taskbar and a desktop with files and folders. Under Windows 95 there is also a consistent "look and feel" across applications. Windows 95 was designed with networking in mind. It has a common user interface for network browsing, resource connecting, and printing. The networking system has also been expanded to provide better connectivity and has become a fundamental part of the operating system. The other major advantages Windows 95 has over Windows 3.1 are additional support for mobile computing (with advances in electronic mail, remote connectivity and pen support), Plug and Play capability (which allows easy setup of hardware devices and laptop computers), and the ability to run 32-bit multithreaded applications.

1.3.2 Hardware Requirements

Windows 95 requires a minimum of 8 megabytes of RAM (Random Access Memory). Increasing the RAM to 16 megabytes will improve performance. The operating system requires about 40 megabytes of hard disk space. Windows 95 requires an 80386 DX or above Intel equivalent microprocessor. These requirements, however, are the minimum. Windows 95 will not perform well under these conditions.

1.4 Windows NT

The Windows NT Operating System was originally designed for distributed computing and networked environments. It has developed into a reliable server platform for client-server applications due to its built-in distributed computing. Windows NT comes in two varieties: a "workstation" version and a "server" version. The workstation version provides a secure, high performance, networked environment for running Windows 32-bit applications. The server version extends this capability by allowing more network connectivity and connected users. Windows NT requires more hardware resources such as memory, disk space and processor power than Windows 95. This is because NT was designed for high performance and robustness. Windows NT also has less backward compatibility so that fewer legacy applications, especially older DOS programs and games, will run on Windows NT rather than on Windows 95.

1.4.1 Architecture

Windows NT has several advantages over Windows 95, specifically: distributed computing, portability, multiprocessing and scalability, and compliance with other operating systems.

Distributed computing is the ability for applications to distribute their work across multiple computer systems. This allows parts of the application to run on hardware which best handles the task to be done. An example would be to run the database portion of an application on an additional server to take advantage of a faster processor or large disk array.

Portability is the ability to run on several platforms and processor types. The Windows NT operating system can run on Intel and DEC Alpha processors, with capability for other platforms in development.

Multiprocessing and scalability refer to the ability to take advantage of several processors running at the same time. Today's computers are being scaled up by combining multiple processors for increased computing capability. NT takes advantage of this design by supporting concurrent tasking of several processors.

Compliance with other operating systems is the ability to run certain POSIX-based applications as well as 16-bit and 32-bit Windows applications and certain OS/2 applications. This allows backward compatibility as well as some portability of applications.

1.4.2 User Interface

The current user interface for Windows NT (version 4.0) is nearly the same as that of Windows 95.

1.4.3 Hardware Requirements

The recommended minimum hardware configuration is 32 megabytes of RAM running on a Pentium processor and a recommended hard disk of 512 megabytes or greater.

1.5 *OS/2 Operating System*

OS/2 is a 32-bit, multithreaded, preemptive multitasking operating system which exploits the capabilities of the Intel 80x86 family of microprocessors. Preemptive multitasking is a more robust way for the operating system to divide system resources. The current version of OS/2, Version 3.0, also known as OS/2 Warp, includes the ability to run legacy DOS and Windows 3.1 applications as well as native OS/2 applications. It also includes TCP/IP and Ethernet support built into the operating system.

1.5.1 Architecture

OS/2 runs applications in what are known as *sessions*. There are four different types of sessions: DOS, Windows, OS/2, and Presentation Manager (PM). There may be any number of these sessions with the exception of the PM session. There is only one PM session, which is similar to the desktop of Windows 95.

DOS *sessions* allow multiple DOS programs to run concurrently. Although DOS is a singletasking operating system, by running multiple DOS sessions several DOS pro-

grams can be running concurrently. Each DOS session is independent and can be configured differently. Both Windows 95 and Windows NT also have this ability.

Windows sessions, called WIN-OS2 sessions, allow the user to run Windows 3.1 applications. Generally, the user will have a single WIN-OS2 session. However, because Windows 3.1 is a cooperative multitasking environment rather than a preemptive multitasking operating system, users may want to run particular Windows programs in their own session. This not only provides the benefit of preemptive multitasking, where the operating system gives processor time to each application in turn, but also provides added crash protection.

OS/2 sessions are text-based console sessions for running OS/2 console applications. They are similar to DOS sessions in looks, but have three important differences. First, OS/2 applications use a flat memory model rather than the DOS segmented memory model. Second, multiple applications can be run concurrently. Third, programs may be run in the background or "detached" from the active session.

The *PM session* is the primary session where the user interacts with the operating system. The PM session allows the user to run Presentation Manager programs (similar in Iook and feel to Windows 3.1 programs) and to create DOS, WIN-OS2 and OS / *2* console sessions.

With the advent of OS/2 2.0 and continuing to OS/2 Warp, a single PM program, called the Workplace Shell (WS), is started when the system starts. This shell sits on top of PM and transforms the OS/2 user interface into a completely object-oriented system.

1.52 User Interface

As described above, the Workplace Shell is a PM application that provides the user with a completely object-oriented user interface. Everything on the OS/2 desktop is an object. The user can obtain a context menu for any object by clicking the right mouse button on the object. The context menu for different objects will be different depending on what the object is or can do.

Applications can create new types of objects and register them with the operating system. When the user sees an object created by the application, e.g., a Describe word processing document, he or she can interact with the object directly without starting the application that created the object. For example, the user might drag the object over to the printer icon and drop it on the printer. The user sees only the printed document-he or she does not need to see the underlying application. In this way, the Workplace Shell provides full objectivity to OS/2.

1.5.3 Hardware Requirements

OS/2 requires a minimum of 4 megabytes of RAM to run, but any OS/2 system should, from a practical standpoint, have at least 8 megabytes. Increasing RAM to 16 megabytes significantly improves performance, but increasing memory beyond 16 megabytes does little to improve performance. The OS/2 operating system requires about 40 megabytes of hard disk space. An additional 40 megabytes should be reserved for the swap file (virtual memory). OS/2 requires an 80386 or above Intel microprocessor.

1.6 UNIX Operating System

UNIX is a very powerful and portable multi-user operating system designed to utilize the computing capability of modern computers and support software development. UNIX is written primarily in the high-level C programming language and makes very few assumptions about the underlying architecture, allowing it to run on a wide variety of computer systems, from a 32-bit microprocessor to a supercomputer. There are two predominant versions of UNIX in commercial use, namely AT&T's System V and the University of California at Berkeley's BSD (Berkeley Software Distribution). However, there are a large variety of other "flavors" of UNIX available for commercial use.

1.6.1 History

UNIX was developed at AT&T's Bell Laboratories in the early 1970's. It was modeled closely after an experimental operating system called Multics and known in its earliest working version as PDP-7. After several revisions, the operating system was rewritten in the C programming language for portability which would eventually lead to its widespread popularity and success. In fact, the C programming language was initially developed for the purpose of revising UNIX into a system-independent operating system. By the mid 1970s, UNIX was in broad use at universities and research facilities such as the University of California at Berkeley, which made several additions to the operating system and released its own version, UNIX BSD. Although no significant industry standard of the UNIX operating system exists, AT&T's System V and UNIX BSD are the most common. The majority of other UNIX flavors emulate one of these two as closely as possible.

1.6.2 User Interface

UNIX is a non-traditional operating system designed primarily for writing, installing, and maintaining large software projects. Unlike prevalent single-user, multitasking operating systems commonly found on modern personal computers, like OS/2 or Microsoft Windows, UNIX runs primarily on workstations. UNIX has no standard graphical user interface (GUI) and utilizes a command Iine interface with the user interface, similar in format to that of Microsoft's DOS. Although this interface is rather cryptic to the newcomer, it provides powerful, interactive access to the services of the operating system. The three main advantages UNIX provides are multi-user access to a single computer, a high level of interaction between the user and the operating system, and convenient data sharing through the file system. These improvements over other operating systems make UNIX an ideal software development environment.

UNIX has traditionally been characterized as a "big system" operating system, despite its portable design. Recent additions of a desktop environment which runs on top of the operating system provide the added benefit of a "user-friendly" graphical user interface. An increasing number of tools readily available for UNIX has led to its common use in the business environment.

1.6.3 Architecture

UNIX can be viewed as a multi-layered operating system, connecting the user to the hardware controls and devices of a computer. At the lowest level, UNIX interfaces with the hardware controls of a system through software known as the *system kernel* or *kernel*. The kernel consists of two main parts, the process control as well as the file management and input/output (I/O). The process control performs memory management, inter-process communication, and process scheduling. The file management portion of the kernel drives peripheral devices and organizes the file subsystem. UNIX performs all of its I/O through reading and writing files, providing a homogeneous interface that handles all device and inter-application communication.

The shell and the compiler components reside in the layer above the kernel. The compiler components are used by programming language compilers during software development. The shell is a command line interface, serving as the user's interface to the operating system and as a means to connect to the compiler components. It consists of a large group of services, functions and links to the underlying operating system controls. Some common shells used with the UNIX operating system include C shell, Bourne shell, Korn shell, and Tenex shell. Many other variants of shells exist as well.

1.6.4 Hardware Requirements

The memory requirements and hard disk space vary greatly depending upon which platform is used. The UNIX system vendor should be able to provide the necessary hardware requirements.

1.7 Conclusion

Many times the operating system you choose will be dictated by the hardware you own or are going to purchase. If you are able to choose an operating system, you must consider what type of applications you wish to run, the hardware requirements, and the anticipated future stability of the system you are considering.

2.0 DATABASE TECHNOLOGY

by DSA Systems, Inc.

Client-server computing is moving into the mainstream of corporate information systems. With this migration comes the need for client-server applications that can access enterprise-wide data. Much of this data is currently stored on large computers running a database tool that enables applications to access the information. At the other extreme are desktop computers containing smaller files of data just as important to an application. Client-server computing provides an open environment where a variety of systems can function together successfully. The techniques used to build, store, re-

trieve and manipulate these databases vary greatly, as do the database application's capabilities and operating requirements.

2.1 High-End Solutions

Figure 2.1 – Multiple client system utilizing client-server database.

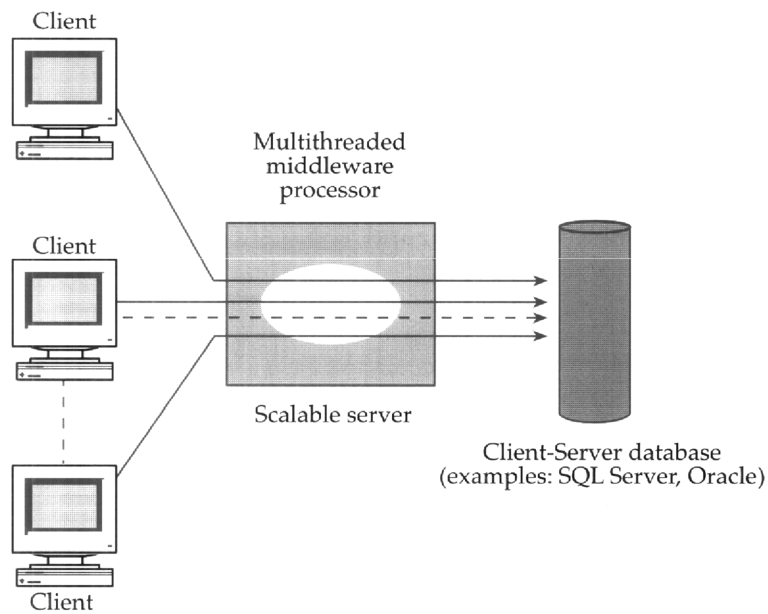

High-end client-server database tools include Microsoft SQL Server, Sybase SQL Server, and Oracle. These tools are designed to handle a large number of transactions and user connections simultaneously. They also provide administration tools to provide feedback and measure performance to determine how well the database tool is performing. These tools usually demand a substantial hardware platform capable of storing gigabytes of information and processing thousands of simultaneous transactions. These high-end tools are designed and optimized with client-server solutions in mind and thus perform well in networked multi-user situations.

2.2 Desktop Solutions

There are also desktop solutions to data storage. These tools are designed to store small to medium amounts of data and can be accessed by small teams of users (1-5 people). Typically, desktop tools provide fewer capabilities and are maintained by the person who inputs and retrieves the data. Examples of such tools are Paradox, Microsoft Access, FoxPro, and dBase, to name just a few. The vast majority of these

database tools use relational database management systems (RDBMS) and flat file system technology.

2.3 Types of Databases

Relational databases use a group of tables which contain attributes to describe the table. These tables are associated with each other by keys (or fields) that are common to each table. Using the key instead of the common data reduces the size of the database and provides connectivity of the data elements. It is this reduction of duplicated data that makes the relational database technology attractive. Flat file databases use a header to define what a record or data element consists of and then provide a list of records sequentially in a file. To access a record, the file is scanned until the data is found. If data is common to multiple records, it is duplicated in each record to maintain the structure of the file. Object oriented (00) databases are new technology that is starting to be used. They provide data storage based on individual instances of data objects and usually provide an 00 front end to access the data. The storage format is usually proprietary, but standards are being developed.

2.4 Database Communications

Database transaction and communication can be handled in several different ways. Open communication standards such as open database connection (ODBC) provide a common programming interface which allows connection to any database that has an ODBC driver associated with it. This open standard is based on structured query language (SQL) which is an ANSI standard for defining and manipulating data. The SQL standard is very popular and widely used. SQL consists of simple, highly flexible commands used to manipulate data stored in tables. It has undergone several changes in the 80s and 90s and will continue to evolve as new database techniques, such as object oriented databases (OODB), become common. SQL is used in almost all relational databases available. However, vendors usually add procedural extensions to the standard in a proprietary fashion to enhance the product. Database providers also usually develop libraries for programming languages so that applications can access the database programmatically. This is also a widely used practice since this usually provides the best performance with the database tool. This, however, limits the program so that it can only access a single vendor's database.

Databases are an excellent (and in some cases an inexpensive) way to organize, manage and retrieve information. The solution you choose should be based upon your hardware requirements, the performance you wish to achieve and the amount of data you wish to maintain.

3.0 NETWORKS

by: Michael Bourke, PhD

3.1 *Introduction*

Advances in clinical applications and activities are being driven by advances in the underlying hardware, systems software, and networks that tie everything together. The major trends in networks have been improvements in the area of speed and reliability, integration of local and wide area networks, growth of wireless, and growth of the Internet for internal use within hospitals.

In order to illustrate these trends, it is useful to compare the state of networking during the mid-1980s with that of the mid-1990s.

In 1985, dial-up modem speeds were 2400 bps and 9600 bps, while dedicated lines ran at 56 Kbps. In 1997, dial-up modem speeds are up to 33.6 Kbps and 56 Kbps, while dedicated lines run at 1.544 Mbps (T1) and 45 Mbps (T3).

A further difference is the growth of ISDN, with a speed of 128 Kbps. In 1985, few departments were on a network. Those that were networked used slow, non-standard systems (e.g., Arcnet running at 1.5 Mbps). In 1997, approximately 90% of all personal computers in large companies are networked. Nearly 90% of these computers use Ethernet (IEEE 802.3).

In 1985, there were no standards for connecting networks. In 1997, TCP/IP has become the universal standard for connectivity.

In 1985, wide area networks (WANs) could only transmit text and numbers. In 1997, WANs are able to transmit voice, images and video at 622 Mbps, and more.

3.2 *Faster Local Area Networks (LANs)*

Traditional LANs use wires to interconnect. Historically, LANs were implemented for small departments and workgroups. As the number of LANs grew, the need to connect them arose, leading to the development of bridges or routers to connect independent groups of LANs. A bridge connects two networks, turning them into a large network. A router does the same thing, but with the advantage of forwarding only the packets whose recipients are outside of the local network.

Finally, the concept of the "backbone network" was developed. The backbone is a single network to which all other networks are attached. These constituent networks are sometimes called "segments". Several segments can be assembled together into a backbone network, using bridges and routers. These LANs are adequate for traditional administrative and clinical data, but they do not have sufficient bandwidth to accommodate new data types such as large images and video. Traditional Ethernet transmits data at 10 Mbps, but this nominal throughput decreases drastically as more users are added to the network, making its effective throughput no more than 5 Mbps. Compare this to the throughput needs of imaging applications, where a single diagnostic-grade X-ray can consume 160,000,000 megabits. In addition, the architecture of LANs, with routers, makes it difficult to accommodate time-sensitive data or to synchronize the arrival of voice, data, and video. LANs were designed to accom-

modate traditional data — the traditional data of databases — "name, date of birth, admission date, comment." On a congested LAN, it is not important for the packets to arrive in a particular order or time. The receiving computer will put all the packets together. However, with voice and video data, the packets must arrive in order and at sufficiently regular intervals to provide an undistorted output. Phone lines were designed to accommodate time-sensitive data and voice. They were certainly not designed to accommodate the volumes of data generated by a video application.

The Internet was not designed to accommodate either voice or video. Rather, its packet switching architecture is that of the store-and-forward model. This new generation of multimedia applications is causing the growth of fast Ethernet technology and various switching techniques. The traditional standards are IEEE 802.3 (Ethernet) and IEEE 802.5 (Token Ring). The bandwidth problems of LANs are being addressed by switching and Fast Ethernet technology.

3.2.1 Fast Ethernet

Fast Ethernet uses the same access control technique as standard Ethernet — carrier sense multiple access with collision detection (CSMA/CD). This means that all stations are free to access the wire at the same time. When two users access the network at precisely the same time, a collision occurs which is detected by software on the network interface card (NIC). Each station waits and then re-transmits. It simply increases the nominal throughput by a factor of 10, but it does nothing to solve the problems of time-sensitive data. Currently available Ethernet NICs support either 10 or 100 Mbps.

3.2.2 Gigabit Ethernet

Ethernet running at slightly higher than one gigabit per second has recently been proposed by a consortium of large, well-known networking vendors called the Gigabit Ethernet Alliance. While it probably won't be widely deployed before 1998, it provides a viable alternative to ATM. The problem with Gigabit Ethernet is that it can not be implemented over the existing Ethernet wiring infrastructure. A fiber optic cable may be required in such a case. Once again, even though Gigabit Ethernet provides a one-hundred fold improvement over lo-megabit Ethernet, it still has the timing problem.

3.2.3 Switched Ethernet

As previously stated, the major weakness of Ethernet-type networks is that they are based on contention. Because of this, much of the potential bandwidth (as high as 40%) is consumed by the collisions. This is solved by the use of "switching". Switching is a commonly used technique for phone systems, When two phones are connected, the phone system reserves a block of bandwidth for communication. The phone system switch has the ability to absorb the bandwidth of all the connecting lines, and it switches the conversation onto a backbone network that has reserved a certain portion of capacity for this conversation. In Ethernet, the user has no guarantees about bandwidth. The technique of switching has become widespread. Figure 3.1

illustrates this concept. The left side of the figure shows a traditional 10 Mbps Ethernet backbone, consisting of wiring hubs and workstations. Each workstation connects to the hub at 10 Mbps. The hub itself has a throughput of only 10 Mbps.

Figure 3.1 – Traditional Ethernet versus Switched Ethernet.

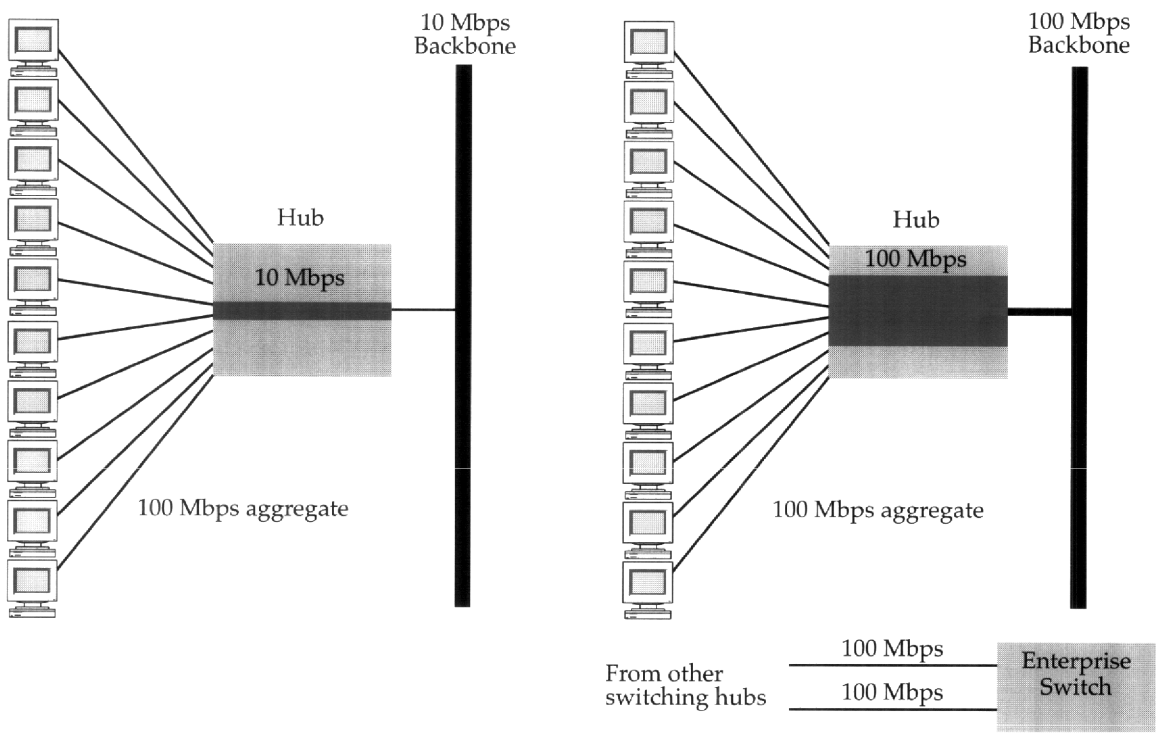

Today's multimedia/graphics applications on a Pentium-class machine can easily fill the capacity of this network. The aggregate bandwidth produced by ten workstations on one hub is 100 Mbps, while the bus of the hub can only accommodate 10 Mbps. While this network appears to be hierarchical, it is completely flat. Everyone has to share the same 10 Mbps. In reality, the bandwidth to any workstation is probably around 250 Kbps. The right side of Figure 3.1 shows the case with switching. The hub has been replaced by a switching hub, with an internal throughput of at least 100 Mbps. The workstations no longer share the same bus inside the switch. The result is that each of the workstations receives a dedicated pipeline of 10 Mbps. There is a high-speed mechanism for switching them onto the backbone. The connection between each workstation and the switch in effect becomes a different segment, with its own 10 Mbps. The device accumulating the workstations onto the backbone is no longer a bottleneck. Similarly, each departmental switch is connected to an enterprise switch by its own 100 Mbps cable. Inside the enterprise switch, messages are switched back and forth over a high-speed backplane, whose speed of 1,000 Mbps is far greater than the aggregate incoming bandwidth of 300 Mbps.

These switches can be installed without the need to replace the existing Ethernet NIC. In addition, the same type of Ethernet packets is used and management principles are similar. This is not the type of packet switching that occurs in a WAN. The system still uses collision detection. It is hoped that these high speeds can accommodate time-sensitive data like voice and video.

3.2.4 Asynchronous Transfer Mode (ATM)

Asynchronous Transfer Mode is the first protocol to be designed for both LANs and WANs. In addition, from the beginning it was designed to handle all types of data — traditional data, voice, graphics, and video. It does this by using a technique called cell-switching. ATM has been defined to run over a variety of media, both fiber optic and coaxial cabling. It has also been designed to run at different speeds. It starts at 25 Mbps, then 51.84 Mbps, and thereafter in 51.84 Mbps increments up to 2.4 Gbps, with the possibility of further upward growth. This means that ATM can be used desktop to desktop across the country, without the need to interface LANs to WANs. However, such end-to-end connectivity requires organizations to discard most of their existing investment in networks. A further obstacle to ATM's use is the rapid growth of 100 Mbps Ethernet and the imminent appearance of 1 Gbps Ethernet. For this reason, ATM has been slow to penetrate the desktop, instead staying in the realm of WANs where it is experiencing strong growth. Eventually, ATM will be installed at every desktop. In the future all PCs will be sold with multimedia and desktop conferencing capability. A new generation of real-time, long-distance multimedia applications will demand ATM throughput to the desktop.

3.3 *Integration of* ***LANs*** *and* ***WANs***

We continue to see a trend to the integration of LANs and WANs. Initially we had standalone PCs and workstations, then they were connected in departmental networks. These departmental networks were then aggregated upward onto a backbone network, linking all departments. But, full connection of LANs over a WAN has eluded us. Similarly, full connection of the home to the work environment has also eluded us (see the discussion of ISDN, ADSL, and cable modem). This is because LAN technology and WAN technology have been antithetical to each other. Each was designed without the other in mind. They differ along many dimensions: protocols, cabling, data types supported, routing methods, capacity, reliability, and cost. The fact that the WAN differed in every aspect from the LAN caused extra expenses, transmission inefficiencies, and management problems (Figure 3.2). The right side of the figure shows the evolution of WANs; the left shows the evolution of LANs. The drawing is divided into stages, each marked by a letter. Empty boxes represent switches, darkened boxes represent routers.

Figure 3.2 – LAN Evolution versus WAN Evolution.

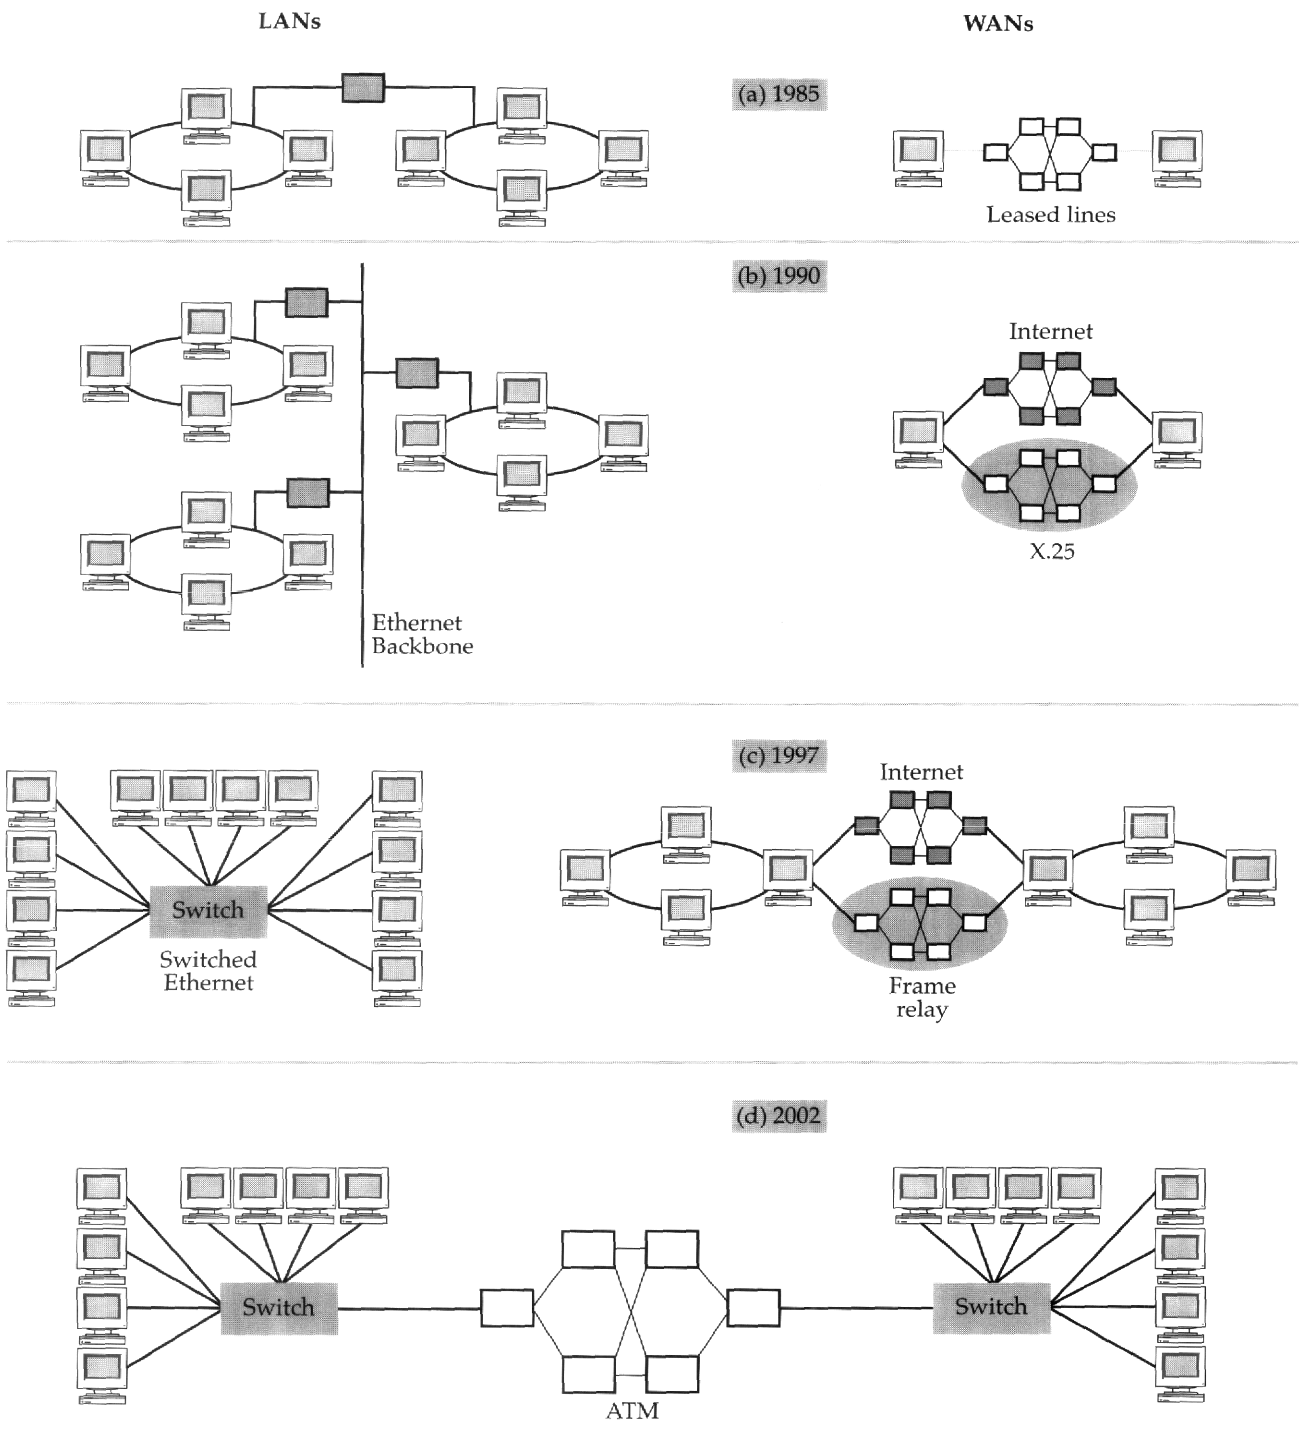

Around 1985, it was possible to connect two computers over AT&T's long-distance network either by dialing or by obtaining a leased line. A dial-up connection would take a different path each time and it was less expensive, but the quality and throughput were lower. The leased line provided the same path for the connection each time. This method was more expensive, but provided higher quality and throughput. WAN connectivity was limited, but, by the same token, LAN connectivity had just started to grow. Most LANs were small and departmentally oriented. There were few demands to connect enterprise LANs across the country.

The period of the late 1980s and early 1990s saw the emergence of large backbone LANs using routers. Now that a large site can be networked, the next step would be to demand connection of these backbones over a WAN. Unfortunately, the two leading WAN technologies for this period, X.25 and TCP/IP, were both unable to support the features and bandwidth needed for connection of backbones. With X.25 and switches, AT&T and other network providers set up a special network that provided the benefits of leased lines, but without the need to lease a dedicated line. Instead, the user would dial into the X.25 network, where his message would be broken down into packets. Each individual packet could take a different route through the network. At the end, the packets would all be reassembled. Because of expense and limited throughput, X.25 was used to connect individual computers or dumb terminals connected to a host. TCP/IP-based Internet used routers and did not provide the type of connection needed to connect backbones coast-to-coast. The Internet, while good for e-mail, could not guarantee the bandwidth needed for transaction processing applications to span multiple backbones.

By 1996, LANs started to use their own switching technology, creating higher LAN throughput and demanding greater WAN bandwidth. Fortunately, new WAN protocols were introduced, which were a big improvement over X.25. These protocols include: frame relay, switched multimegabit data services, and ATM (to be discussed in following sections). It was now possible to connect not just individual computers, but all computers over a WAN. At the same time, the Internet was trying to overcome the limitations of routers and a protocol that could not accommodate time-sensitive data. The Internet Engineering Task Force, the standards body for the Internet, is in the process of defining a way to specify quality of service, which would result in guaranteed bandwidth.

Ultimately, probably by 2002, LANs and WANs will have completed their convergence whereby they are connected through ATM from end to end. Routers and switches will have merged their functionality in devices that perform switching, but also have routing tables. The TCP/IP protocol will accommodate time-sensitive data. Multimedia applications like the electronic medical record will run over the WAN transparently.

3.3.1 Asynchronous Transfer Mode (ATM) for WANs

ATM is being used more and more to integrate LANs across a WAN. The major telecommunications carriers (e.g., AT&T, MCI, Sprint) are installing it. The vendors who provide the backbones of the Internet (e.g., UUNet) are installing it.

An example of the use of ATM for the transport of multimedia data on a WAN is a project called Starbright World.' In this project, ATM technology is used to provide seriously ill children with a three-dimensional multi-user environment that allows them to meet, communicate, and play with children in other hospitals. The system consists of specially configured Pentium PCs, named PC Pals, with color monitors, microphones, speakers, and video eyes. It transmits both voice and full-motion video. It even provides alternate modes of input for children who are too sick to speak or type. The workstations will accept a range of input devices, including keyboard, joystick, mouse, and other more specialized devices to accommodate disabilities.

The participating hospitals include: Mount Sinai Hospital, New York, NY; Boston Children's Hospital, Boston, Massachusetts; Lucile Salter Packard Children's Hospital, Palo Alto, California; Children's Hospital of Pittsburgh, Pennsylvania; Children's National Medical Center, Washington, DC; Children's Medical Center of Dallas, Texas; and University of California at Los Angeles Children's Hospital, Los Angeles, California. Each hospital has a 10 Mbps Ethernet network which connects to an ATM switch, which in turn connects to a 45 Mbps commercial WAN. In addition, each hospital runs ATM all the way to desktop computers.

3.3.2 Integrated Services Digital Network (ISDN)

In order to provide full LAN connectivity across a wide geographical area, including connections to the home, one must overcome issues of speed, cost, and reliability. A modem connection from a home-based PC to a clinical information system is not fast or reliable enough to support all the performance characteristics and functionality customary on a full LAN (e.g., fast database access). If a faster communication line is desired, one would have to pay much more. ISDN is an effective compromise. It is fast (128 Kbps), inexpensive, and does not require a permanent connection.

ISDN takes advantage of the existing copper wire that the telephone company has installed running to each home. There are different levels of ISDN, but the most frequently used is 128 Kbps, which begins to approach the throughput of some LANs. In addition, ISDN is defined to work with the TCP/IP protocol, and provides universal connectivity. With the appearance of ISDN and the development of the point-to-point protocol (PPP, part of the TCP/IP suite), it became feasible to connect homes, offices, and branches to LANs through dial-up lines, thereby making it possible to participate in LAN services (e-mail, database access, exchange of graphics data, and acceptable exchange of video data) at a reasonable cost.

The University of California at San Francisco Medical Center (UCSF) has installed an ISDN network to support telemedicine from physicians' homes. This network allows physicians to query the hospital's diagnostic imaging database so that they can perform remote diagnosis of patients and complete paperwork. This has resulted in increased productivity of physicians whose homes are connected to this network.

The UCSF Medical Center also uses ISDN to connect three local health clinics which have LANs. The 56 Kbps connections the clinics used previously were too slow, while upgrading to a T-l line was prohibitively expensive. In the future, UCSF Medical Center plans to use its ISDN network to support video conferencing to allow doctors to remotely view echocardiograms and monitor patients from their homes.'

3.3.3 Asymmetrical Digital Subscriber Line (ADSL)

While ISDN has existed for 15 years, asymmetrical digital subscriber line (ADSL) is a recent arrival. Like ISDN, ADSL makes use of the existing telephone wiring in homes. It provides communication at two speeds (hence asymmetrical). Users can upload to the network at a speed of about 600 Kbps, and can download at a speed of about 6,000,000 Kbps, or 6 Mbps. The rationale for this asymmetry is the fact that most end nodes (physician office, home, mobile user) spend most of their time downloading from a central source. In an environment where clinicians need access to image data, this service can prove very useful. ADSL is not widely deployed; however, several local telephone companies are testing this technology in a number of markets. While ADSL may seem to be a good alternative, it is uncertain whether it can really run at 6 Mbps on the older household wiring. Moreover, ADSL creates a permanent connection, which has implications for IP addresses. This means that a permanent network address must be assigned to each ADSL device. II' addresses are already in short supply. Despite these problems, ADSL could be a major threat to the growth of ISDN, which is still troubled by problems of inconsistency across providers. In addition, ISDN will always have less bandwidth than ADSL (only 2% of ADSL).

ADSL has been tested in Canada, where the Royal College of Physicians and Surgeons of Canada, the Ottawa Civic Hospital, and the University of Ottawa conducted a trial of ADSL for use in physician education. A course on emergency medicine was created for physicians who could not attend the campus. From their offices, doctors used personal computers to work with simulated emergency medicine cases. These cases contained a mixture of full-motion video, audio, animation, text, and clinical images. The physicians were able to order electrocardiograms, ultrasound, and X-rays of the simulated patients, as well as use tutorials and reference materials.'

3.3.4 Cable Modem

Another promising network technology is that of cable modem. Cable television is in 60% of American homes, and nearly 100% of hospitals. Cable television uses the same kind of coaxial cabling as used by LANs, with the capability of throughput as high as 30 Mbps. However, there are two big differences. First, television signals are analog waveforms, while LAN signals are digital. Second, television messages are only one-way, from the broadcaster to the receiving home. Some hospitals, like hotels, have some limited two-way capability, permitting the room to communicate with the network. The first difference is overcome by a device called a cable modem, which is really a router that would send the digital signals to the PC instead of the television. The second difference will require some changes to the distribution equipment of cable vendors, by which they would be able to accept a greater volume of upload data. Currently, some cable modem products have the upload performed over a telephone line and not the cable. Additional issues to be resolved concern the origin of the broadcast, segmentation of the network, and routing of messages.

While throughput could potentially reach 30 Mbps, it is anticipated that initial transmission rates will be in the 64 Kbps to 512 Kbps range. While tests of cable modem are being conducted in many large cities, there are almost no standards for the

cable modem. In order to address this lack of standards, a large number of well-known vendors (television manufacturers, PC manufacturers, and networking companies) have formed an ad hoc standards organization called The Broadband PC Council. Unfortunately, there is a parallel effort called the Broadband Link Team also working on competing standards.

3.3.5 Frame Relay

ISDN, ADSL, and cable modem are techniques for connecting a small office, home, or mobile employee to the corporate network. They are not suited for connecting large entities. This latter is done by frame relay. In the mid-1990s frame relay became the most common method of connecting LANs across a WAN. Unlike the store-and-forward approach of the Internet, frame relay is for real-time connections, and allows real-time data dialogs,

Frame relay is similar to an older protocol called X.25 but is much more efficient and less expensive. This is because frame relay uses less error checking than X.25. Since X.25 does not assume the network to be reliable, it performs error checking at every node. Frame relay, on the other hand, assumes an underlying reliable network and thus does not have to do error checking at every individual node. In addition, frame relay sends packets more efficiently. X.25 uses a technique called time-division multiplexing. Each incoming line is given a slice of time, regardless of whether it has anything to send. Frame relay, on the other hand, uses statistical multiplexing, which assigns line packets to any unused slots making it a cheaper alternative. Frame relay runs over T1 and T3 lines and provides a migration path to ATM.

3.3.6 Switched Multi-Megabit Data Services (SMDS)

Switched multi-megabit data services (SMDS) was introduced around the same time as frame relay. It is a high speed (45 Mbps) LAN to WAN technology, which can carry voice, data and video, and offers a migration path to ATM. Moreover, SMDS does not require a permanently wired setup but uses a phone dial-up type of connection instead. Unfortunately, SMDS has been slow to catch on and will probably never enjoy widespread deployment. Most companies now choose either frame relay or ATM. Today, only one major long-distance carrier provides SMDS, and its growth trails both frame relay and ATM.

3.4 Wireless Network

An important component for both the LAN and the WAN is the wireless network. This is an important development for hospitals. The computing platform should allow the clinician the flexibility to access patient care data while on the move within the hospital. Wire-based LANs, in contrast, do not permit this. Wireless networks can be based on either radio wave or microwave communication. With microwave, the stations have to be in line-of-sight, while radio waves can travel through walls and around corners. One issue with radio wave communication is security. Anyone with a radio receiver can intercept the transmissions. This can be solved by encrypting the

message before transmission. Another issue is speed; so far, it has proven difficult to achieve speeds greater than 2-3 Mbps. Compare this with 10 Mbps Ethernet, which is rapidly giving way to 100 Mbps.

There are four distinct aspects of wireless networks:
1. Wireless connections to LANs,
2. Wireless connections to WANs,
3. Wireless bridges between LAN segments, and
4. Wireless connection of peripherals to desktop devices.

3.4.1 Wireless LANs

There are two parts to a wireless LAN solution: the wireless client adapter and the access point. The wireless adapter is usually a PC card (formerly known as PCMCIA) adapter. The term "wireless LAN" is a misnomer, because in most cases wireless LANs do not replace wired LANs. Instead, they create wireless extensions to wired LANs. Wireless LANs need an "access point" which is a stationary device that attaches to the wired LAN (Figure 3.3). An antenna links the wireless clients to the wired LAN via the access point. Wireless LANs have significantly less bandwidth than wired LANs — only 1-2 Mbps compared to 10-100 Mbps. Moreover, interference from other electrical equipment limits wireless coverage and throughput.

Figure 3.3 – Roaming wireless stations attached to cabled network.

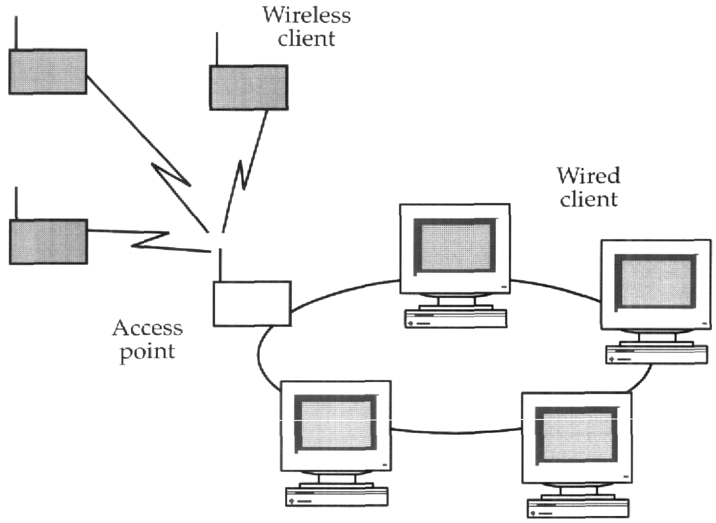

Several wireless technologies are currently available for the LAN including infrared communication and two types of spread-spectrum radio frequency, DSSS (direct sequencing spread spectrum) and FHSS (frequency hopping spread spectrum). DSSS products are the fastest, but FHSS offers better security because FHSS transmitters are

constantly hopping from one frequency to another in no particular pattern. DSSS may be easier to crack, since its signals can be deciphered by determining the spreading code. However, wireless LAN products have additional security features, including data encryption, data scrambling, and user IDs.

FHSS and DSSS operate over the radio airwaves in the ISM (industrial, scientific, and medical) bands of the electromagnetic spectrum (902-928 MHz range and 2.4-2.484 GHz range). Infrared operates between the visible part of the electromagnetic spectrum and the shortest microwaves. The IEEE committee has included both FHSS and DSSS techniques in its draft of the 802.11 wireless specifications.

A wireless network is employed by Saint Joseph's Hospital in Denver, CO, to provide clinicians with all pertinent information concerning a patient, as well as the hospital's entire knowledge base, which is contained on two 14-drive CDROM stacks. St. Joseph's selected laptop computers (over hand-held units) using PC cards to attach to its 10 Mbps Ethernet network. This system is based on 2.4 GHz frequency-hopping spread spectrum and operates at 1.6 Mbps. The units have an operating distance of up to 1,000 feet in open spaces.'

3.4.2 Wireless WANs

With the growth of national systems for paging devices and cellular telephones, supported by the growth of hand-held computers, it was natural that a national wireless network for computers would be proposed. This is called cellular digital packet data (CDPD). In 1992 a consortium of cellular carriers united to define a uniform standard for transmitting data over existing cellular telephone channels. Computers are able to use the existing cellular telephone infrastructure because the cellular channel is frequently idle, even in the middle of conversations (as much as 30% idle time). Using the IP protocol, the CDPD will allow mobile computers to connect to traditional WANs.

3.4.3 Wireless LAN Bridges

In some situations, it is necessary to install a wireless bridge to overcome physical obstacles that stand in the path of network cabling. This occurs in many campus situations, where a man-made structure (street) or natural hazard (river) interrupts LAN cable runs. One solution would be to divert the network through the local telephone company, but this can be expensive. Instead it is possible to install microwave transmitters on either side of the obstruction. Such wireless bridges offer speeds of up to 2 Mbps, which is faster than the 1.544 Mbps of the T1 WAN transmission standard at distances up to 25 miles.

3.4.4 Wireless Connection of Peripherals

Peripherals like mice and printers can now be connected to their home-base desktop computer through an infrared connection. A further use of this infrared technology is to synchronize a laptop computer with a desktop computer. Physicians about to perform rounds could download the rounding schedule to a laptop or handheld computer. Then while making rounds, updates can be made to the calendar which, upon

returning to the office, could be uploaded into the physicians' desktop or laptop computers. These systems generally have a maximum useful range of 4 to 6 feet and an effective throughput from 50 Kbps to about 100 Kbps.

3.5 *Other Issues*

Some LAN technologies, like Token Ring and FDDI, allocate bandwidth more effectively than Ethernet. Recall that Ethernet allows any user to access the network at any time. Simultaneous access results in collisions, which are resolved by the station's backing off and retransmitting. Ethernet is called a "non-deterministic" network because it is impossible to predict how much bandwidth one will obtain, thus making it impossible to predict response time. Token Ring and FDDI, on the other hand, are deterministic. That is, they do not use contention to allocate bandwidth to users. Instead, they circulate an electronic signal, called a token. If a user does not have this token, he/she may not use the network. Such networks make much more effective use of bandwidth — up to 90% and higher. However, they are being pushed into the background by fast Ethernet, switched Ethernet and ATM.

3.5.1 Token Ring

Token Ring, an IEEE 802.5 standard, was promulgated by IBM in the mid-1980s, and was predicted to surpass the number of Ethernet installations. Token Ring, however, was never adopted on a wide scale. The prices of Token Ring NICs were always higher than those of Ethernet. There were shortages of Token Ring chips; and Ethernet was closely associated with the openness of the Internet. Token Ring has a speed of 16 Mbps. Because of its deterministic protocol, about 14-15 Mbps of that is usable. Many large IBM users are currently replacing their Token Rings with Ethernet.

3.5.2 Fiber Distributed Data Interface (FDDI)

Fiber distributed data interface (FDDI) is a very stable and powerful technology. FDDI has a bandwidth of 100 Mbps and has been an IEEE standard since the mid-1980s. Its fiber optic cable is impervious to noise, has a redundant ring built into it for fault tolerance, and it is deterministic, guaranteeing each workstation a predetermined amount of bandwidth. FDDI has been used as a departmental LAN when high bandwidth to each desktop is needed (for example, medical imaging applications). It has also been used as a backbone network, because it utilizes bandwidth much more efficiently than 100 Mbps Ethernet (over 90% efficiency, compared to Ethernet's <u>maximum</u> of 40-60%). While FDDI is still being selected as a high-speed backbone, there is no doubt that it has been hurt by the arrival of Switched Ethernet and ATM. In many cases, FDDI is being used as a temporary solution until ATM standards firm up and ATM prices come down, Furthermore, while an FDDI ring can stretch for 20 miles, it is still not meant to be a long-distance WAN.

3.5.3 Virtual Networks

The first LANs were constructed as large, flat entities. As more users were added, performance degraded (see comments on Switched Ethernet). To get around this, LANs were partitioned into 10 Mbps segments using routers. Routers examine incoming packets to determine the best route, convert network protocols, and filter packets so they do not flood the network. Routers do a lot of work so they become bottlenecks; moreover, router management was quite complex. In these flat networks, users were only associated with other users on the same physical segment. With the appearance of switched LANs, a new kind of segmentation became possible — logical segmentation. The ability to segment a network logically is called a Virtual LAN. This makes it possible to manage groups of people who are scattered over multiple physical segments of the network. Broadcasts and multicasts only go to the members of the virtual LAN — they do not flood each segment, which might happen to have one member of the given work group. This is illustrated in Figure 3.4. Here we see a network with an enterprise switch acting as the backbone for six physical segments, managed by switches. There is a switched network for six areas: the Oncology Department, the Lab, Medical Imaging, the Medical Building adjacent to the hospital, Surgery, and the Medical School several blocks away.

These six physical segments can be divided logically into any number of virtual networks to support a variety of clinical teams. For example, the Breast Cancer Team (solid dot) consists of seven members who are scattered throughout the six physical segments. The Colon Cancer Team (open dot) consists of 8 members, who are also scattered throughout the hospital's campus. It is possible to configure and mange each team virtually, ensuring that team broadcasts do not go to all the nodes of all the segments on which team members are located. In addition, the reconfiguration of the member's locations is simplified by the management software provided in many vendors' systems.

3.6 *Internets and Intranets*

Between 1994 and 1996, the Internet experienced an explosive growth. Several factors contributed to this:

1. A decision was made to allow commercial activities over the Internet,
2. The Internet backbones started to be upgraded to fiber optic cable,
3. New tools for using the resources of the Internet (browsers) and for developing network aware applications (Java) appeared.

All these factors permit the development of a new generation of applications which can take advantage of Internet resources just as easily as they can access their own database. In addition, this new spurt in Internet exposure has further strengthened the status of Internet protocols as <u>the</u> universal standard. This has also caused TCP/IP to undermine the hegemony of Novell's IPX/SPX. Any vendor of LAN products must now accommodate both protocols.

Figure 3.4 – Two virtual networks within one large switched network.

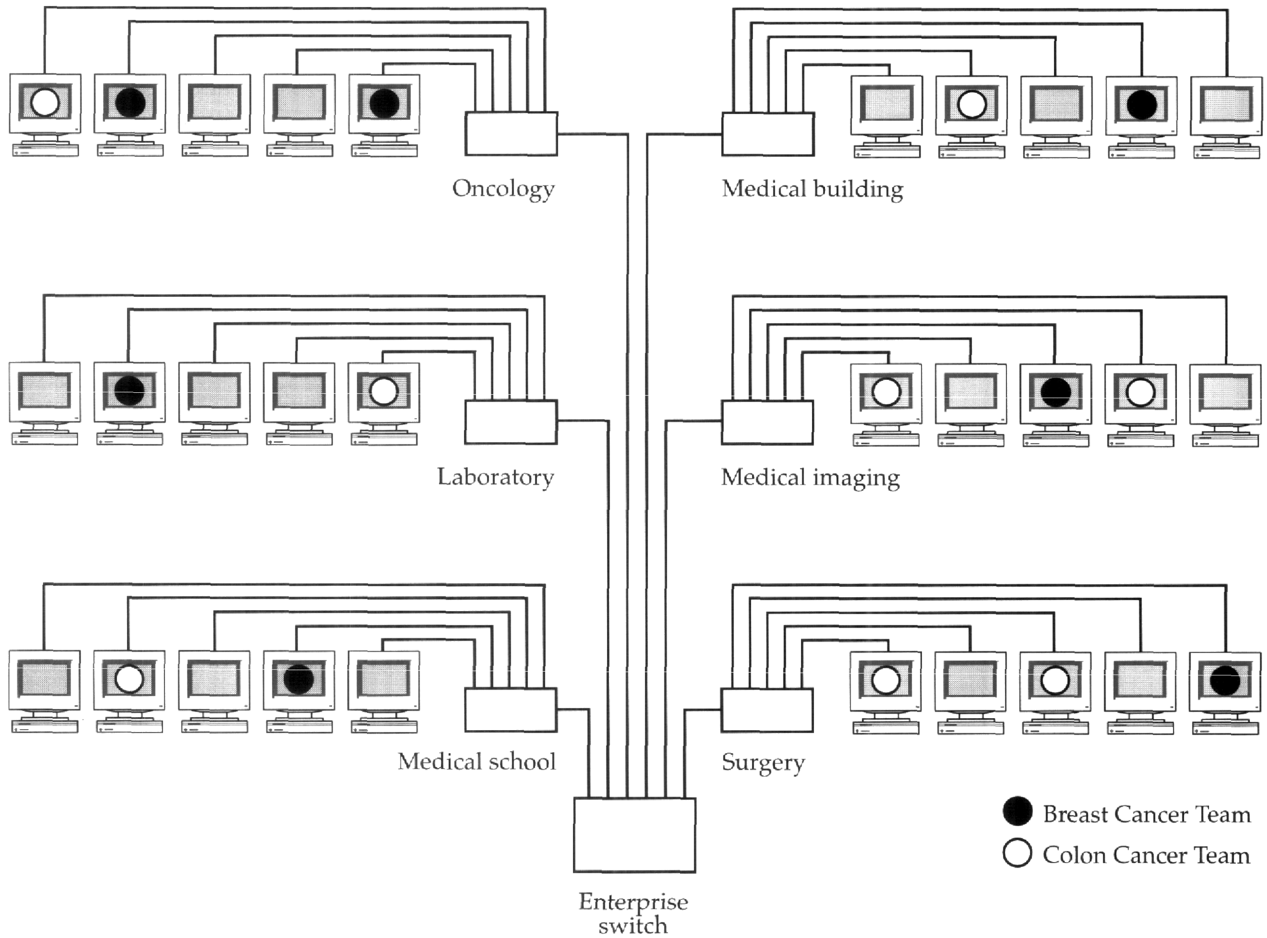

3.6.1 Intranets

Given the spurt in Internet use, the next logical step was to apply this new set of Internet tools and standards to networking situations that are internal to an organization. This is called "intranetworking." In an interesting irony, the key to internal systems integration has been found in the external environment. Intranets can be either LAN-based or WAN-based. Intranets address long-standing problems facing system integrators, namely:

1. Lack of a cross-platform client interface,
2. Lack of cross-platform development tools,
3. Lack of network-aware applications, and
4. Lack of cross-platform database access tools.

All this changed with the arrival of Web servers, browsers, Java, and gateways to back-end databases. The concept is shown in Figure 3.5. The browser (universal interface on the workstation) accesses an HTML page, which is network-aware. When the user clicks on a hyper-text section of the home page, he can be transported across the network to another home page. He can also trigger a network-aware application, written in Java. In addition, when he reaches another home page, he can obtain access to data resources through new data gateways.

Figure 3.5 – Intranet architecture.

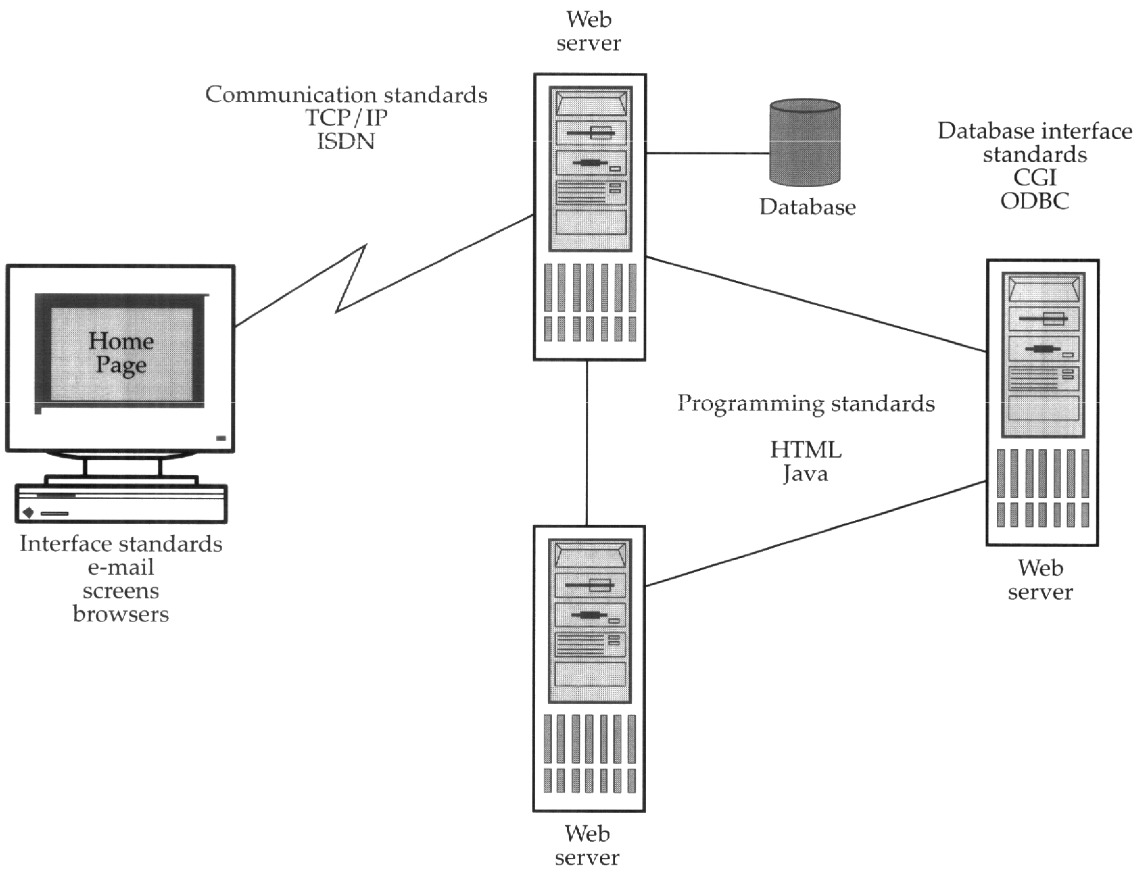

The initial uses of hospital intranets focus mainly on administrative activities like distributing human resource data, submitting expense forms, electronic publishing of paper-based documentation, and information searches. Over time, however, hospital intranets will be used to support the electronic medical record, allowing each ancillary system to publish the data that it has on the in-house patients, including dictation, images, and video. The intranet will also be used to support integration through HL7, functioning as the vehicle for the exchange of HL7 EDI transactions.

Promima Health Systems is testing an intranet to link its nine hospitals and eliminate paper documentation. Eventually, it hopes to have 200 to 300 users taking advantage of this intranet. Constructed using workflow software and Internet-enabled office automation tools, Promima's intranet provides access to data such as on-line policies and procedure manuals, travel vouchers and accounting information, Promima is also considering additional applications such as phone books, events calendars, continuing education courses, and a home page for each of its 300 physicians.[5]

The Hollings Cancer Center developed an intranet application to reduce the confusion and paperwork involved in clinical trials by medical researchers. Protocols can range in size from 25 pages to hundreds of pages. Previously, the protocols were distributed in several stages, resulting in the delivery of hundreds of copies to physicians, nurses, administrators, and data collection personnel. The volume of paper was exacerbated by the fact that protocols always undergo multiple revisions. Putting each protocol on a Web server solves the paperwork problem by maintaining a single, up-to-date copy. This intranet also helps community physicians stay aware of new clinical trials, allowing them to refer patients to new trials by filling out an intranet form.'

Becton Dickinson, a large healthcare products company, provides real-time data to its 250 medical sales representatives in the field. The company equipped its sales force with laptop computers to browse the company's web page to view and update information on hospitals, doctors, and distributors. Representatives also use web-based forms to enter contracts, sales call plans, as well as contact and product data. When representatives want to bring up a list of hospitals, for example, they click on the hospital hyperlink on the home page. This sends a request to a SQL database; the results of the query are formatted as HTML and downloaded to the browser. The downloaded hospital name is itself a hyperlink, and the representative can drill down into further details by clicking on the hospital name.[7]

3.7 Conclusions

The foregoing discussion of networking technologies, together with examples of network use by healthcare providers, illustrates the inexorable trend to integrate all components of healthcare organizations, regardless of distance (LANs and WANs can be integrated), regardless of organizational function (administrators, researchers, hospital-based caregivers, and community-based caregivers), regardless of mobility (in the office, in the field, or on rounds), and regardless of information type and quantity (text data, voice data, image data, or video data). This convergence of networking technologies not only presents formidable challenges for the managers of information systems, it also creates even more fundamental problems for the managers of the healthcare entities themselves, who will have to redesign their organizations in order to make effective use of these integrated computer networks.

3.8 References

1. Schurr A. Wish upon a Starbright: a dose of technology speeds recovery of ailing children. PC Week. 1995;12(48):N1-2.

2. Schurr A. ISDN lets MDs make long-distance diagnoses. PC *Week*. 1996;13(1): N3.

3. Ottawa doctors offered multimedia training. *Newsbytes*. 199-L

4. Clark E. The wireless ward: Saint Joseph Hospital brings information technology to the bedside. *LAN Magazine*. 1995;10(13):143.

5. Rice V. Health group gets intranet treatment. PC *Week*. 1996;13(18):51.

6. Carr J. Intranets deliver cross-platform access to company data via dedicated networks. *InfoWorld*. 1996;18(8):61.

7. Youngworth P. Mobilized forces. *Data Based Advisor*. 1996;14(9):42.

4.0 WORKSTATIONS AND COMPUTERS

by: Michael Bourke, PhD

There are two major trends affecting workstations and computers. The first of these is the issue of scalability — the ability to expand or shrink your computer depending on the needs of the business. The second issue is that of multimedia — the ability of a computer to process a variety of data modalities — images, moving images, waveforms, and sound). The multimedia trend is a result of the recent advances in processor technology.

4.1 Multimedia

A decade ago, the term "multimedia presentation" meant a slide show accompanied by a sound track. Slides consisted of text and graphics on transparencies, while the sound track provided an overlay of commentary or music. On computers, graphics usually had to be imported in a clumsy way into the text-based application. Dedicated machines were often used for this purpose.

4.1.1 Chips for Multimedia

With the availability of powerful CPUs and high-capacity disk drives in the mid 1990s, it became possible to build applications that had tightly integrated text, graphics, and sound capabilities. Of course, for years medical imaging equipment has had specialized chips for converting the various imaging modalities into digital signals. Now these chips are being used in consumer products, making it possible to construct a multimedia workstation out of a personal computer.

4.1.1.1 Digital Signal Processors

A class of chips, called digital signal processors (DSP), has been developed for digitally encoding and decoding various data modalities — analog voiceband signals, stereo audio, fax-modem transmission, and video signals. In addition to encoding and decoding, DSPs can manipulate data, compressing it for faster transmission and better storage, detecting patterns in the signal, and filtering out noise (extraneous signals) for applications such as fax-modem line conditioning, digital telephone

answering, voice messaging, simple speech recognition, voice annotation for word processing, and spreadsheets. When digitized, voiceband signals, stereo audio, fax-modem transmissions, and video signals can be manipulated and compressed by a DSP to ensure high-capacity storage, faster transmission rates, reduced background noise, greater fidelity, and wider dynamic range. DSPs perform these special operations 10 to 1,000 times faster than general-purpose microprocessors.

There are two types of DSPs: general-purpose and special-purpose. While general-purpose DSPs can be used for multimedia and stereo sound reproduction, video manipulation and compression need the higher performance of special-purpose DSPs. While some general-purpose DSPs are still used for satellite photo processing and medical image manipulation, it is usually the high-end special-purpose DSPs that are used in medical imaging equipment. This includes magnetic resonance imaging (MRI), computed tomography (CT) scanners, positron emission tomography (PET), and ultrasound.

While most of the industry standards for video processing are fairly well defined, the standards for audio processing are currently undergoing revisions,

4.1 .1.2 CPUs for Multimedia

Many vendors are changing the architecture of their CPU chips to accommodate efficient processing of multimedia data, in particular, video. Sun Microsystems and Hewlett Packard have had this feature in their chips since 1994. The most significant event however, was the appearance of Intel's Pentium chip with multimedia extensions (MMX) in late 1996. The MMX Pentium is up to eight times more efficient for running multimedia applications than the previous versions of the Pentium chip.

4.1 .1 .3 Media Processors

Other vendors such as IBM are working on media processors which are auxiliary DSP chips that specialize in multimedia functions such as moving picture expert group (MPEG) encoding and decoding, videoconferencing, graphics, sound, and telephony features. Media processors will be used in older Pentium machines to offload multimedia processing, and perhaps in some MMX Pentiums since, despite the MMX extensions, it is not clear how many audio, video, and communications functions can really be performed by the main CPU working alone. For example, the decoding of MPEG-2 video at 30 frames per second could consume all the processing power of a 233 MHz MMX Pentium with MMX, which does not count the audio tracks, which also consume huge amounts of MIPS.

4.1 .1 .4 Multimedia Workstation

We will soon see multimedia workstations to support physicians, For example, the Hospital of the University of Pennsylvania, Philadelphia, PA has designed and implemented an image and report workstation which can display multimodal images, provide radiology reports (in both voice and transcribed formats), and perform teleconferencing. The workstation was implemented on a standard UNIX platform with Motif as its front end.'

4.1.2 Compression Standards for Multimedia

In order to integrate video into computer programs, it is necessary to have a standard which specifies how the video data is to be digitally encoded and how it is to be transferred internally from CD to the CPU.

The most common compression techniques for various applications are MPEG for video, JPEG for images, and global system mobile (GSM) for audio and voice. While GSM is used in Europe, it has not gained wide acceptance in the United States. At the present, systems must use uncompressed audio .wav files.

One problem with the JPEG compression technique is that it is prone to data loss. Although JPEG compression yields 10:1 to 15:1 compression ratios, there is no guarantee that the image would maintain its fidelity. Telemedicine requires a "lossless" transmission of data. A data compression technique that produces a lossless compression of image data is the wavelet technique, which is being adopted as a standard by the American College of Radiology. Teleradiology software has already been developed using wavelet compression. This software can handle a variety of modalities: X-rays, CT scans, MRIs and angiograms. American Telemedicine International is using this new technology to transmit radiological images from Saudi Arabia to Boston.'

4.1.3 Sound

4.1.3.1 Computer Telephony Integration

Computer telephony integration (CTI) is a technology for interfacing a PBX to a computer network. Historically, each PBX vendor developed its own interface for interacting with a computer network. This made applications expensive, clumsy and technically risky. In recent years, however, PBX makers have opened up their systems, incorporating features of generic PCs, and PC vendors have agreed upon a standard interface for PC-PBX interaction called TAPI.

The open features introduced by PBX makers include using non-proprietary operating systems like UNIX; offloading of administration, console interface, and control from the PBX's CPU to separate PCs; and changing their bus structure to accommodate Ethernet NICs and TCP/IP.

On the PC side, two vendor conventions have been placed in the public domain, greatly contributing to the opening up of PBXs. These telephony interfaces are TAPI from Microsoft and TSAPI from Novell. These application program interfaces work with Windows-based programs to make it possible to place and transfer phone calls, make conference calls, perform automatic faxing after phone conversations, display pop-up screens of customer records as a call comes in, and gather information from customers while they are on hold. This is illustrated in Figure 4.1.

Figure 4.1 - Using computer telephony to integrate databases with phone calls.

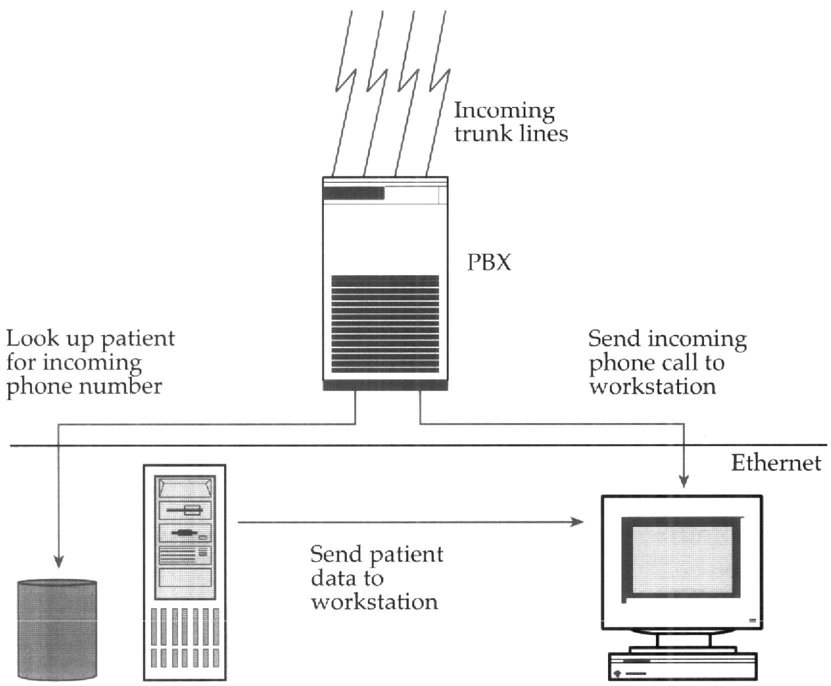

Network servers are displaying more and more telephone features. Small companies are even using telephony-enabled LANs instead of proprietary PBXs. This ties in with the increasing demand to use the Internet for voice messages. Eventually, the PBX will merge with the LAN server.

Another CTI application is voice recognition. One standard is the speech recognition API (SRAPI), backed by Novell, IBM, and others. It runs on multiple platforms, including the Mac OS, OS/2, and Windows. SRAPI lets applications integrate speech recognition with speech synthesis, reading e-mail aloud to the user over the telephone.

CTI is used by the Humana chain, which has a large customer service center of over 600 employees handling customer service calls. With their computer linked to their phones, these employees do advertising, set appointments for their field people, process new subscribers, handle inquiries, and perform other 800-number functions. When Humana customers call in, their data is displayed on the screen simultaneously.[3]

4.1.3.2 Sound Compression

Audio data can be encoded and decoded using a standard called .WAV files. A .WAV file can consume from 1 MB to 10 MB per minute of recorded sound, depending on the quality desired and whether the recording is mono or stereo sound. The size of

these files shows that there is a need for a standard way of compressing audio data — videoconferencing applications make it essential to have a standard for the compression of voice data streams. Unfortunately, the standards for compressing .WAV files are not in place.

4.1.4 Virtual Reality Devices

Advances in digital signal processors lead naturally to virtual reality (VR) applications in medicine, where complex, interrelated forms of information need to be accessed by teams of specialists, often in a real-time situation. The move toward virtual reality in medicine comes from the success of telemedicine, where videoconferencing has brought remote physicians to the bedside, conferring with the patient or local physicians to resolve medical problems. The next logical step is to model virtual environments, which can be accessed over networks like the World Wide Web.

A prototype VR clinical system has been constructed by the University of Washington in its Human Interface Technology Laboratory. They have modeled the emergency room of a local hospital, allowing medical students to get "hands on" experience without leaving the classroom. The wall and equipment of the emergency room were modeled by taking an extensive set of photographs and merging them onto a cylindrical virtual surface. This spatial metaphor is used not only to depict the patient, but also to integrate all the relevant data, For example, instead of clipping X-rays on a wall-mounted light table, a physician can move an X-ray window over the patient's body and display any X-ray data available for that area. Other patient information such as EKG charts or CAT scans can be positioned as virtual objects that float in space in front of the user. The data elements can be moved around and positioned like objects, so that all participants in the virtual operating theater have instant access to relevant data.'

4.2 *Scale Down / Scale Up*

The personal computer started a revolution. Since the PC represents a standard computing platform, it is natural that users would want to use that platform to build bigger applications, culminating in applications that support the whole enterprise. As a result, the once-humble server is increasingly showing characteristics of mainframe computers. Conversely, it is natural that users would want to take this platform and miniaturize it for mobile computing. It is also natural that they would want to extend it to other equipment that processes data in the general sense of the term - medical monitoring devices, PBXs, etc.

4.2.1 Computers That Scale Up

Hardware components (CPU, bus, disks, peripherals, memory) have become standardized. Operating systems have also become standardized (UNIX, DOS, Windows). Granted, UNIX is somewhat fragmented and Windows is controlled by a single vendor, Microsoft, but this is a far cry from the environment of minicomputers and mainframes of a decade ago. In that world every vendor had proprietary hardware, a proprietary operating system, and a proprietary network. If you had an appli-

cation that was slowing down because of business growth (increased transactions, increased users, increased sites, etc.) you faced difficult choices. You could wait until the vendor released a more powerful CPU, or you could move the application to another vendor's platform. Neither alternative was attractive. This situation is similar in the PC world.

Take the example of a hospital that uses a LAN-based order entry and charging system. As the number of uses grows, the CPU becomes overloaded. Figure 4.2 illustrates this point. With a single CPU, when the number of users of an application exceeds a certain level, the computer can no longer process the transactions. User requests are placed in lengthy queues, and all work slows down.

Figure 4.2 - Growth of delays as users are added to system.

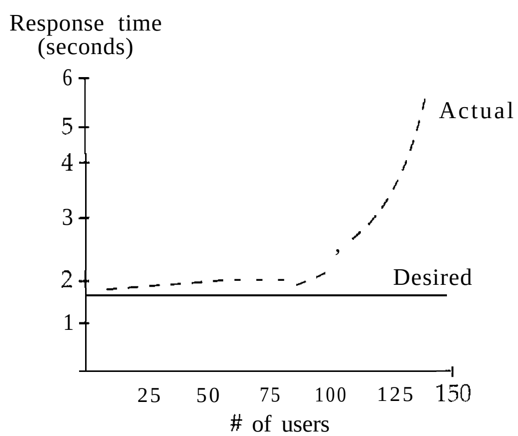

4.2.1 .1 Multiprocessing

One solution is to employ multiple CPUs in a computer. Another name for this is symmetric multiprocessing (SMP). In SMP, all CPUs share the same bus and memory. It would be nice to add on "generic" CPUs in the Intel x86 series, and have a constant linear increase in performance. Windows NT, OS/2, and UNIX can all support the use of multiple CPUs. Unfortunately, at around 32 CPUs the machine runs out of power — the PC bus cannot accommodate the extra traffic resulting from having multiple CPUs which must communicate with each other. Symmetrical multiprocessors consisting of between 2 and 32 CPUs can be used for s and transaction processing, but beyond this limit of 32 CPUs, other technologies must be employed.

SMP with generic PC chips is part of a more general trend: the movement to make LAN servers more like the architecture of mainframes. This is not surprising, since more and more mission-critical, high-performance applications are being moved to LANs. Mainframe features include: 1) multiple CPUs; 2) offloading certain functions to special-purpose CPUs; 3) redundant components like buses, controllers, disks and power supplies; and 4) backups that can be done on the fly, when databases are open.

A further technique for managing server overload is to dedicate different servers to different specialized tasks. One could have an image server, a communications server (for modems and fax), or a printer server. This is illustrated in Figure 4.3.

Figure 4.3 - Dedicated servers.

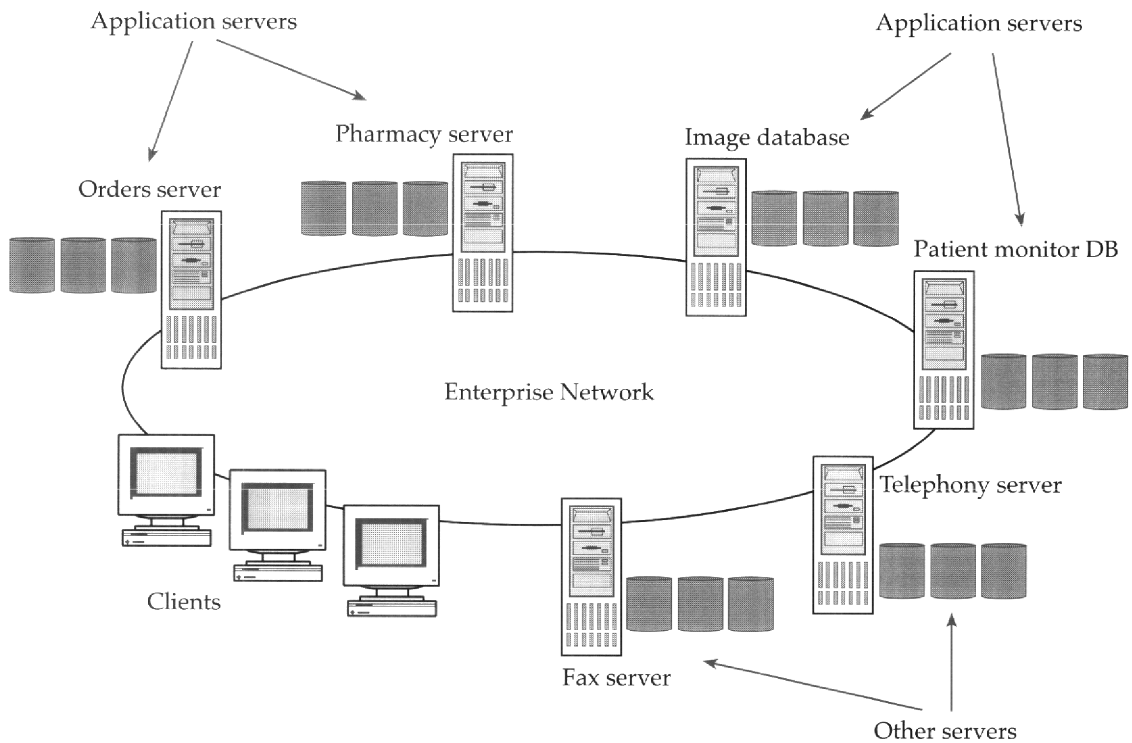

4.2.1.2 Clustering

Clustering is one technique to overcome the memory bus bottleneck to increase performance. For example, one can connect two 16-processor SMP servers, or "nodes", creating one 32-CPU configuration. This is usually cheaper than building one large 32-CPU system. This type of clustering was invented by Digital Equipment Corporation in the late 1970s with its VAX architecture. With clustering, you employ several different boxes, each of which may have multiple CPUs, that are linked with a network connection which increases redundancy.

Clustering can be expanded so that the total number of clustered CPUs far exceeds the CPU limit of SMP machines. There are, however, problems with clustering. In theory, users can get significant performance gains with each additional node. But each node has its own memory pool and system bus, which means programming has

to be done to coordinate the communication of information and sharing of data across the various nodes in the cluster. Thus, performance improvements are not linear, and the bandwidth of the interconnect between the nodes themselves is often too low to sustain high-performance gains. Some vendors offer specialized interconnect technology to overcome this bottleneck.

Another benefit of clustering is that it creates high-availability and redundant systems. Two or more nodes can be tied together in one cluster so that if one node fails, another node takes over automatically, For example, Hermann Hospital in Houston has installed a cluster of two 4-CPU machines to run a patient records system for Hermann's more than 1,400 affiliated providers. In addition to fast response times, this cluster provides the required high availability, disaster tolerance and fail over capabilities.[5]

4.2.1.3 Massively Parallel Processing

Massively Parallel Processing (MPP) offers the greatest possibility for near-linear performance gains. In an MPP configuration, hundreds, if not thousands, of processors can be tied together. Unlike clustering or SMP, each processor in MPP has its own memory and bus and is capable of handling its own tasks while it communicates with the other processors via specialized interconnects and switches. This overcomes the interconnect and bus bandwidth bottlenecks of the other architectures.

There are two drawbacks to MPP: cost and ease of use. A typical MPP system can start at $500,000, compared with $50,000 - $100,000 for an entry-level SMP system. In addition, MPP is more complex to manage than the other architectures since the applications and the database have to be modified extensively to run optimally on MPP systems. This frequently requires that the applications have to be "parallelized" so specific tasks are allotted to specific processors or sets of processors.

4.2.2 Devices That Scale Down

It is also desirable to take the desktop functionality along with you as you perform your job.

4.2.2.1 Hand-held Devices

Hand-held computers are also called personal digital assistants or PDAs. PDAs can recognize handwritten notes and can transmit them into a database. With a stylus, the clinician can check boxes, or can print in commands to the system. PDAs communicate with a LAN by radio waves. Brigham and Women's Hospital in Boston is using PDAs to support its residents. These PDAs contain medical textbooks and drug references.[6]

There are other hand-held devices, like cellular phones and pagers, which eventually will merge with PDAs. The ideal PDA must then be able to do networking, paging, faxing, phoning, scheduling, storage of phone numbers, voice mail handling and note taking — all this while coordinating the details with a desktop computer. Right now, this is difficult and expensive to achieve in a PDA so the universal deployment of PDAs will have to wait.

While PDAs are still best suited for organizing, the latest models now run a wider variety of business applications than before — from word processing and spreadsheets to e-mail messaging, voice memos, and Web browsing. Some new devices even run a "pocket" Windows 95 operating system called Microsoft Windows CE.

Some technical considerations for PDAs:

- **Weight.** PDAs usually weigh less than a pound, work for up to two months on AAA batteries, and are pocket-sized, which makes them handy travel companions for checking appointments and messages.
- **Power.** The AAA batteries used by most PDAs have lives that range from 15 to 100 hours. However, vendor estimates of battery life are based on different assumptions concerning the type of work done on the PDA. For example, a PC card modem on a PDA will drain the battery in a short time, so it is advisable to get an optional AC adapter.
- **Screens.** While screens may vary in size, they all use monochrome LCD technology, which does not make for optimal legibility, even when they're equipped with backlighting. Backlighting must be used judiciously, since it is a power hog.
- **Input.** Most PDAs come with a miniature QWERTY keyboard, but also include a stylus which can be used as a mouse by tapping the touch-sensitive screen. Some PDAs have handwriting recognition software, allowing the user to write on their screens with a pen which the PDA then converts into editable text. Many users find this feature difficult to employ.
- **Speed.** Since PDAs do not have a hard disk they store data in RAM. This makes loading and saving data files fast. The boot up process is instantaneous.
- **Software.** The basic software suite for PDAs usually includes an address book, calendar, notepad and a calculator. Some units — the Windows CE devices for example — provide a light version of Microsoft Office: word processors, spreadsheets, and databases. Since PDAs have to fit all this functionality into just a few megabytes of ROM, these applications have reduced features. If the PDA is used in conjunction with a desktop PIM (personal information manager) such as Lotus Organizer, it is important to have software which synchronizes the PDA appointments and address book entries with those of the desktop unit. Bidirectional reconciliation is important — both the PDA and the desktop should reflect the changes, no matter where they are made.
- **Prices.** Prices range from $300 to $700. This price does not include features like data synchronization, e-mail, extra memory, or various PC card peripherals. In order to do more than simple notes and calendaring, one must spend at least $1,000.

4.2.2.2 Network Computer

The term "network computer" refers to a slimmed-down computer connected to a network, with the network providing much of the storage and processing that has traditionally belonged to the desktop PC. Over the years, desktop PCs have grown in power and complexity. They have added tape backup units, enormous hard drives, fax cards, whole suites of software, and other components. The reason for this was two-fold: 1) tools, data, and devices either did not exist on the mainframe or they

were too expensive or difficult to use; and 2) mainframes had been replaced by more user-friendly servers but the networks were not fast or reliable enough for the user to turn over his storage and processing to the server. The term “fat client” is used to refer to this practice of loading up a workstation with extra hard drives, CD-ROM drives, and application software needed by the user. In contrast, the network computer represents the “thin client.”

Historically, fat clients have created several problems: 1) they are more expensive than regular PCs or terminals; and 2) they created many support and management problems. The first issue is really a red herring. The total cost of ownership (TCO) per year of a $2,500 machine can be anywhere between $10,000 and $15,000 when one factors in support, training, and management of the configuration. It is in TCO where vendors claim the network computer is superior.

Now, with the advent of faster networks and new types of software (downloadable “applets”), many experts are calling for a return to the model of storing most programs centrally. There are several alternative solutions to the network computer. Some vendors propose using a powerful terminal, while others recommend a sealed PC. Some store the programs on a server; some have programs in ROM; others store the programs on the sealed hard drive.

Some vendors claim that the network computer will sell for around $500. But, in reality, the price will probably be closer to $1,000.

There is much confusion with regard to the definition of TCO. While a network computer’s TCO might be 40% lower than that of a fat client, this scenario assumes that the fat clients are poorly managed. If the fat clients are well managed, then the difference in TCO may be as small as 20% or less. Do you want to change your computing paradigm for that?

Some hospitals, like the North Arundel Health System in Baltimore, Maryland are successfully using thin clients. The hospital has about 250 Windows-based terminals for accessing network databases.[7]

4.3 Conclusion

Eventually, all technical innovations lead to discussion of business processes. It is clear that computing is now a continuous process which can integrate an enterprise — there is no activity too small to use a computer; there is no activity too large to use a computer; there is no data beyond the reach of digitization. However, all the business processes of healthcare were designed with the assumption of discontinuities in computing and data. Consequently, there are major discontinuities in our clinical and administrative processes which must be remedied before we can achieve the maximum benefit from the advances in chips and computers.

4.4 References

1. Kohn D. Imaging World Shows DICOM Standard Works. *Health Management Technology.* 1995;16(4):20-23.

2. Ryan M. Distance Health Care is Latest Medicine. Electronic *Engineering Times.* 1996;899:55-56.

3. Jordan P. The Customer for Computer Telephony. *VARiness.* 1996;12(1):7A.

4. Brown C, Dose of VR Prescribed for Medical Data. *Electronic Engineering Times.* 1996;194:34

5. Smith L. Prescribe the Right Hardware to Fit Your Organization's Needs. *Health Management Technology.* 1996;17(11):28-31.

6. Freeman DH. Boston's Wired Hospital. *Forbes.* 1994;153(8):S114.

7. Russell J. Amdahl Thins its Focus. *Communications Week.* 1997:657:8-9.

5.0 INPUT/OUTPUT DEVICES

by: Michael Bourke, PhD

Input and output devices are used to digitally capture both administrative and clinical data, or to provide this digital data later on in one form or another. Unfortunately, traditional information technology has prevented the digital capture of much data produced by healthcare activities. Even where data capture has occurred, for the most part the data has been fragmented and non-integrated. The problem was that the process of data capture (or its presentation) was intrusive — it interfered with the business process — one had to step away from the patient to view the data or sit down at a keyboard to enter the data. This has changed dramatically in recent years. We have gone from a limited capability (characters, numbers, and text), to digitization of images, and now to digitization and presentation of any data modality imaginable (voice, body movements), without intruding on the activity itself. The input and the output of data are becoming indistinguishable from the clinical activity and the medical instruments. The very performance of the activity triggers the digital recording of data (sensors, speech recognition). Flat display panels are unobtrusive. Images and video can be retrieved in real time, and soon will be displayed one the lenses of one's glasses.

5.1 Cards

The term "cards" refers to a class of input/output devices whose size has gone from that of a shoe box in the 1980s to a credit card in the 1990s. They are a great example of the microminiaturization trend -- the shrinking of the electronic circuits in computer chips. There are many types of cards, some are used to hold patient data, others are used to hold electronic money or encrypted passwords. Others are used as peripheral devices — hard drives, modems, etc.

5.1.1 PC Cards

The original term for this technology was PCMCIA; however, it has been shortened to "PC card". There are several classes of PC card which differ in profile and thickness. They are all, however, about the size of a thick credit card. Because of their weight and size, they are ideal for use in laptop computers to provide the functionality of a fax, a modem, a parallel printer connection, a network adapter card, a pager, cellular phone, etc. They weigh several ounces and consume almost no power. Some of them

have multiple functions — modem and network connector. In other cases, the user has to pop out one PC card and pop in another.

In 1996 St. Luke's Episcopal Hospital in Houston, Texas, installed laptop PCs equipped with PC cards for wireless networking. The first department to benefit was the hospital's respiratory therapists who use this wireless network to schedule treatments, perform documentation, and capture billing data. This system is also connected in real-time to the main HIS, allowing instantaneous access to changes in orders and patient status.[1]

5.1.2 Smart Cards

These are credit-card size devices that patients carry. They can store varying amounts of data, ranging from a patient's insurance data to the patient's medical record. While such cards are very convenient, they have two drawbacks: they are easily lost, and their data can become out of date easily. They can be swiped across a reader, like credit cards. Alternatively, they can be inserted into a PC card slot to be read.

The Department of Defense has recently started to deploy a battlefield medical information system that uses smart cards. In a combat exercise, troops landing on a beach were linked to the information systems on a hospital ship. The forward surgical companies near the beach used wireless communications to link to the ship offshore. Each soldier carried a smart card with basic medical information about themselves. Field medics used a rugged hand-held computer to send smart card data to the shipboard medical information system. The hand-held systems also downloaded current data onto the smart cards, to provide up-to-date information about the soldier's current condition and treatment.[2]

5.2 External Busses

A computer "bus" is the set of wires connecting all the computer's components so that they can get data to and from the CPU. Like the traditional city bus of the 1950s, everybody gets on the same bus to "go downtown". Devices that plug into the internal bus include internal disk drives, internal modems, external modems, printers, scanners, etc. This was a very clumsy and error-filled process. The user had to open up the computer, insert the device's controller into a slot, close the computer, then plug the device into the port on the controller.

The goal of an "external bus" is to be able to connect a variety of external devices to the workstation in a common manner, and to do so without having to open up the equipment. Instead of having to fit the connector to the correct port on the back of the workstation, an external bus is built into every device, running through them all. This allows for easy daisy-chaining of devices. A further goal is high throughput. One of the problems with existing printer and modem ports is that they were designed in an era when networks were slow or non-existent, and output devices were slow and character-based. Using a bus permits the easy in-line attachment of such diverse devices as modems, printers, disk drives, CD-ROM, and tape backup units.

The external bus also represents a solution to the docking station issue. As is well known, portable computers like PDAs and laptops have weight, power, and size restrictions. Vendors have created docking stations (proprietary extensions of the com-

puter), into which the computer is plugged upon returning to the office. External busses provide the functionality of docking stations and will be cheaper and more flexible because they are not proprietary. It is very possible that clusters of these external busses will be distributed throughout the hospital for use by clinicians on the fly, instead of requiring that they return to their office.

5.2.1 Universal Serial Bus (USB)

Universal serial bus, or USB, has been an emerging standard for a while; however, vendor fighting over the final standard has prevented its wide acceptance. It will probably become widespread by the end of 1997. The universal serial bus (USB), whose development committee included all the major names in the PC industry, is a low-speed version of the Firewire (see below), with speeds up to 12 Mbps. USB allows users to daisy-chain up to 127 plug-and-play peripherals. It has a 5-volt power line for these peripherals, so that they can be added and removed without powering down or reconfiguring the computer. The USB will be used to simplify and standardize the interconnection of printers, mice, keyboards, digital joysticks, telephones and telephone networks, scanners and modems.

5.2.2 Firewire

The proper designation for the firewire is IEEE 1394, a high-speed serial bus developed by Apple originally for high-bandwidth links to the PC for digital-videodisk players, cameras, and camcorders. IEEE 1394 supports 63 devices with a throughput of 100 - 400 Mbps. IEEE 1394 differs from USB in that it is designed to meet the high-speed requirements of peripherals that deal with high-bandwidth time-sensitive data, such as monitors, video adapters, and CDROM drives. Like the USB, IEEE 1394 has been slow to gain acceptance, with many vendors taking a “wait-and-see” approach. Clearly, however, this is a much needed technology in light of the advances in videoconferencing, and eventually will become a standard.

5.2.3 SCSI

SCSI stands for small computer system interface. The SCSI standard was developed during the early 1980s and allowed up to 7 devices to be daisy-chained to a single SCSI controller. It allows connection of disk drives, backup units, and CD-ROM without having to resolve device addresses and find free space inside the computer. The use of SCSI has proliferated recently, with the trend to assemble huge disk arrays (see RAID below) for the creation of s. However, SCSI has given way to Ultra-SCSI, which has throughput of up to 40 megabytes per second. But SCSI standards will probably be replaced by IEEE 1394 or a technology called Fiber Channel - Arbitrated Loop (FC-AL), which has been proposed to ANSI by a consortium of vendors to ANSI. FC-AL has throughput up to 200 Mbps, with fault tolerance — well suited for today’s multimedia applications.

5.3 *External Storage*

5.3.1 RAID

Redundant arrays of inexpensive disks (RAID) is a technology for creating subsystems of disk drives that have high capacity, redundancy, and fault tolerance. There are five different levels of RAID, each with a different set of features. RAID 1, for example, specifies disk mirroring. In disk mirroring there is a primary and a secondary disk drive. All the data written to the primary drive is also copied to the secondary drive. The highest level, RAID 5, can provide both fault tolerance and redundancy with as few as two drives. This growth in disk-array technology is being matched by growth in arrays of tape drives to back up these disks when the files are still open and being used by applications.

St. Elizabeth Medical Center in Covington, Kentucky, has installed a RAID-5 configuration to support its patient accounting and medical records modules. This configuration runs on two SMP servers and has 80 gigabytes of storage.[3]

5.3.2 Data Warehouses

Data warehouses are very large databases which integrate the diverse data resources of the hospital for purposes of research and analysis. Frequently, they store from 5 to 10 years of historical data. When this data is not summarized, but is kept at the level of the transaction, such databases can reach 100 gigabytes and more. They must accommodate efficient consolidation of the data along different views or dimensions (payer, diagnosis, time period, procedure, medicine, cost, etc.). This requires extraordinary amounts of processing power and new ways of physically structuring the database. Traditional relational databases, with tables consisting of rows and columns, do not provide effective storage and access. In addition, traditional uniprocessor computer architectures do not have the power to efficiently handle complex queries that require the processing and combination of gigabytes at the same time. Frequently, machines with up to 8 CPUs are used. Sometimes several machines, each with multiple CPUs, will be tied together in a technique that is called "massively parallel" processing. One solution is to develop a "star" architecture for the database, using a large table in the middle for patient demographic data, with spokes of summarized data radiating out from this center. Each spoke contains arrayed data for a particular dimension of the patient: diagnosis, stay, treatment, drug, etc.

The term for this type of activity is on-line analytic processing or OLAP. This term is in contrast to on-line transaction processing (OLTP), the activity that inputs the basic transactions of the organization — registrations, transfers, orders, services, charges, bills, etc. No single clinical database has the capacity to support these two antithetical types of activity — OLTP and OLAP. A hospital will need different types of databases. While relational DBMSs have evolved to become very efficient for OLTP, where a given transaction is restricted to only a handful of rows in several tables, they cannot accommodate OLAP, where a single transaction, such as a query, can require processing of every row of dozens of tables for 10 years back.

Intermountain Health Care, headquartered in Salt Lake City, Utah, has constructed a which will ultimately serve 12,000 to 15,000 users. To do this, Intermountain chose a massively parallel configuration from IBM. The data in the warehouse will come from such diverse sources as hospitals, insurance companies, and outpatient clinics, and it will support both clinical and financial decision making.”

5.3.3 Object-Oriented Databases

There is increasing progress toward creating the all-electronic medical record. This requires digitization of all patient data — medical images, doctors’ dictation, handwriting, video, biomedical instrumentation data, etc. While these diverse data modalities can be effectively encoded and manipulated with the new generation of digital signal processors, they cannot be accommodated by traditional relational database management systems which require data that fits neatly into tables (eliminating voice, images, video, and waveforms). The solution is to use Object-Oriented (00) databases. In addition, 00 databases are very efficient in processing large, complex transactions, like checking out multiple volumes of a patient’s medical record and holding on to them for several days. This is in contrast to a more atomistic transaction like registering a patient, which can be handled easily with a traditional relational database management system.

Scott & White Memorial Hospital in Temple, Texas, recently installed an object-oriented DBMS that combines traditional text data with images and digitized lab data.’ By combining data that has traditionally been stored in different systems, and by displaying this data on a single screen, the hospital can provide physicians with a more integrated view of a patient’s condition and history.

5.4 Displays

5.4.1 Thin Screen

Traditional CRTs take up a lot of room on the desktop. A CRT whose screen measures 13” on the diagonal might be up 15 inches deep, and the depth increases as the width increases. A further problem with the CRT is its sensitivity to electromagnet radiation. In areas like MRI, where magnetic fields are high, LCDs (liquid crystal displays) are preferred.

Some vendors are developing thin LCD screens, for desktop computers, which are 10 to 22 inches in width. The smaller units are just one-inch thick. While this technology is expensive for the larger units, prices are expected to drop dramatically as sales volumes increase. Such thin screens are being employed in the Long Beach Memorial Medical Center in California, where physicians and nurses use flat-panel PCs equipped with 10.4 inch touch-sensitive displays that fit flush against the building’s walls. The PCs can thus be located wherever they are needed, whether for emergency or operating room functions, lab or pharmacy work, or patient records. They will be tied to the medical center’s backbone network to ensure integration of data.[5]

In order to achieve a wider screen, but maintain a thin profile, one has to use plasma display technology — a mixture of helium, neon and xenon gases which are sandwiched between two substrates. For example, a 42-inch plasma-display panel

can cost as much as $12,000, but it will have a definition (see HDTV below) suitable for diagnostic activities.

5.4.2 HDTV

HDTV stands for high definition television, The original specification, developed by the television and cable industries, called for a video format of 60-frames per second using a particular scanning technique (interlaced) with a pixel format of 1,920 x 1,080. The TV signals used were digital, not analog as in current cable TV. A resolution of 1,920 x 1,080 pixels permits the use of HDTV for diagnostic activities. As this standard was nearing adoption, however, a consortium of PC manufacturers interjected itself into the process, opposing the interlaced scanning and suggesting a different scanning method, different screen resolutions, and a variety of speeds (24, 30, and 60 screens per second). This argument resulted in the promulgation of a much broader standard, now called DTV (digital TV). DTV permits a variety of specs which will be subject to interpretation by hardware and software manufacturers. Equipment developed to this new standard is expected late 1997 - early 1998.

5.5 *Voice*

5.5.1 Speech Recognition

In speech recognition, sound waves are digitized. Using artificial intelligence techniques, this digitized data is then analyzed by programs which determine the word, phrase, and sentence boundaries. Then the programs parse the sentence grammatically and check the lexical meanings. All of this is very resource intensive. It uses enormous amounts of disk space to store the sounds, and enormous amounts of CPU cycles to process them. Speech recognition applications must also accommodate background noise, the physical condition of the speaker (hoarse, an accent), and new words. Vocabularies can vary from 20,000 words (niche application like Radiology) up to 60,000 words (for genera1 diagnostic work).

Speech recognition systems have been resisted by physicians, primarily because of lack of technology to process continuous speech — causing the physician to pause after every word or phrase. Thus they could not support the physician's real-time work flow and style. More recent products now support real-time continuous speech, prompting with speech synthesis when they find an unfamiliar or ambiguous item. In addition, these newer systems can also recognize voice commands requesting the retrieval of patient demographic information. A prime candidate for such technology is the radiologist. Historically, radiologists have dictated patient information using a tape or digital recording system. The recording was then transmitted to a transcriptionist for keying, then returned to the radiologist for review and signature, which was a lengthy and wasteful process. One vendor, working with radiologists from Memorial Sloan-Kettering Cancer Center, New York, and the Massachusetts General Hospital, Boston, has developed a system that lets users dictate, edit and electronically sign their own reports in real time using a personal computer. The benefits are faster report turn-around, lower transcription expenses, and potentially improved

confidentiality. Such productivity improvements are important for tertiary care facilities which can generate hundreds of thousands of radiology reports per year.[6]

5.5.2 Speech Synthesis

Speech synthesis refers to the conversion of on-line text to spoken words. Another term for this is “talking database.” Speech synthesis is more reliable than speech recognition because the computer doesn’t have to interpret anything. It simply converts the ASCII text of a computer file into synthetic speech and plays it back over speakers. The more advanced speech synthesis systems employ an expert system that helps the computer pronounce items that do not comply with phonetic rules: foreign words, abbreviations, acronyms, and symbols. Speech synthesis is a very stable and easy-to-implement technology, which is actually underutilized. The reason for this is that it probably works best in situations where it is paired with the less developed and more expensive speech recognition (e.g., in-patient care, where hands and eyes are focused on the patient).

5.6 *Exotic - Virtual Reality*

The technologies discussed thus far are familiar and conventional. However, advances in virtual reality have produced other input and output devices which will eventually find use in healthcare, given the fact that healthcare providers need to have their hands and eyes free.

5.6.1 Wearable Computers

While they are not found in the consumer market, wearable computers are being used in industry for real-time, mobile, and data-intense applications. One vendor has produced a lightweight cap onto which can be mounted small units for speech recognition, speech synthesis, video, graphics, and text. It consists of a small display, a microphone, and a speaker which can be attached to different styles of these caps.[7]

A significant event is the development of computer displays that can fit into or onto one’s glasses. While pixel density is rather low right now, it can be expected to grow to VGA density and higher in the near future.

5.6.2 Cameras

Everyone is familiar with the technique of mounting a small camera on a surgeon’s head to broadcast an operation to a remote auditorium. This was a forerunner to the surge in videoconferencing that has occurred in the past several years. The camera is the input device. The interesting aspect of videoconferencing is that it has multiple simultaneous data streams - voice, video, and traditional data. This requires the interaction and synchronization of different input/output devices (speaker, microphone, CRT, disk drives, etc.).

The ITU (International Telecommunications Union) has set a number of standards for videoconferencing. The H.320 standard addresses videoconferencing over ISDN and other circuit switched networks. The H.323 standard, finalized in 1996, is an extension of H.320 and is designed for

local area networks (LAN), wide area networks (WAN) and corporate Intranets that are generally packet-switched networks. The H.324 standard covers traditional telephones.

5.6.3 3-D Digitizers

Some vendors are manufacturing 3-D digitizers which can scan three-dimensional objects, such as the human body, into the computer. These objects can then be rendered (the process of applying different colors and textures) and converted into formats that the common graphics programs can read. An example of this is “Mr. Jack” from the University of Pennsylvania. Mr. Jack is an anthropomorphic data set representing the human figure. It has 71 segments, 70 joints, 136 degrees of freedom and a 17-vertebra spine. Mr. Jack is accompanied by special software for the control and positioning of this complexly jointed object, such as placing a virtual dummy into a vehicle.”

5.6.4 Sensors

Everyone is familiar with the sensors that are attached to the patient to capture an array of physiological data (heart rate, temperature, respiration rate, etc.). Initially, this data was displayed in analog form on the patient monitor. An early advance was the digitization of the waveforms, which also permitted storage of a greater variety of physiological data which can then be linked with diagnosis and medication data in clinical databases.

Similarly, vendors are providing technologies that will display data coming out of the instruments used to treat a patient. Technologies such as CAT, PET, MRI and X-ray, until now, have been two- and three-dimensional static displays of patient data, but they can now be turned into dynamic displays. For example, the Johns Hopkins Hospital has installed a tracking device onto a scalpel so that a surgeon can see the relative position of a scalpel within a 3D MRI rendering of the patient’s brain.’

But there are other types of sensors, developed by the virtual reality industry, such as a set of wires that can be lightly attached to the face to sense the constriction of facial muscles, which then can be fed into a computer to guide the animation of a cartoon face.

5.7 Conclusion

This chapter has tried to show the advances in information technology that have made data capture and presentation nearly universal and unobtrusive, embedded in the activities themselves, rather than requiring separate activities, which detract from the basic activities of healthcare. In closing, I would like to raise a non-technical issue. Historically, all the clinical and administrative processes of healthcare were set up with the assumption that data had to be captured and presented through additional processes, that it was not integrated, and was incomplete. As a result, a convoluted set of work flows, functional responsibilities, paperwork, and organizational procedures were created. Consequently, in order to obtain the maximum benefits of these new input/output devices, equal attention should be paid to re-engineering the activities that they impact.

5.8 References

1. Carr J. Wireless network market. VARiness. 1997;13(3):127-131.
2. Bass B. Hospital system survives first battlefield trial. Federal Computer Week. 1995;9(8):1.
3. Sperling E. Optika goes vertical with Medipower. Computer Reseller News. 1996;678:22-23.
4. Appleton E. The right server for your Datamation. 1995;41(5):56-58.
5. Stevens L. UniSQL solves the paper chase. Open Computing. 1995.
6. Johnston SJ and Needle D. The PC still matters. InformationWeek, 1996;582:48-50.
7. Darnell T. Devices as dynamic as the people who use them. Health Management Technology. 1997;18(2):78-82.
8. Brown C. Wearable computing looks good. Electronic Engineering Times. 1997; 947:153-154.
9. Park WT. VR is dead, long live VR. Digital Media, 1993;3(1):25-27.
10. Toigo JW. Healthcare: biomedicine meets Buzz Lightyear. Digital Age. 1996;15(8):20-22.

6.0 GLOSSARY

Compiler Components – Used by programming language compilers during software development.

Compliance (with other operating systems) – The ability to run certain POSIX-based applications as well as 16-bit and 32-bit Windows applications and certain OS/2 applications.

Context **Menu** – A menu of options specific to an object accessed by clicking once on the object with the right mouse button.

Desktop Environment – Runs on top of the operating system with a "user-friendly" graphical user interface.

Detached Programs – Programs that may be run in the background from the active session.

Distributed Computing -The ability for applications to distribute their work across multiple computer systems.

File Management (of kernel) – Drives the peripheral devices and organizes the file subsystem.

Multiprocessing – The ability to take advantage of several processors running at the same time.

Multitasking – Allows the operating system to execute several programs simultaneously.

Multithreading – Allows an application to divide tasks into smaller executable sections.

Operating Sysiem -A program that controls the execution of application programs and acts as an interface between the end user and the computer hardware.

Plug and Play – Easy setup of hardware devices and laptop computers.

PM **Session** – The primary session where the user interacts with the operating system.

Portability – The ability to run on several platforms and processor types.

Preemptive Multitasking – A more robust way for operating systems to divide system resources.

Process Control – Performs memory management, inter-process communication and process scheduling within the kernel.

RAM – Random Access Memory.

Scalability – The computers today are being scaled up by combining multiple processors for increased computing capability.

Server -Allows network connectivity and connected users.

Sessions (DOS) – Allows multiple DOS programs to run concurrently.

Sessions (Windows) – Called "WIN-OS2 Sessions", allows the user to run Windows 3.1 applications.

Sessions (OS / 2) – An application running in OS / 2.

Shell – A command line interface, serving as the user's interface to the operating system and as a means to connect to the compiler components.

Swap File – Virtual memory of the hard disk.

System **Kernel** -Hardware controls of a system through software.

Workplace Shell (WS) – Sits on top of the PM and transforms the OS/2 user interface into a completely object-oriented system.

Workstation – A secure, high performance networked environment for running Windows 32-bit applications.

INDEX

Spacelabs Medical, Inc.
15220 NE 40th Street, P.O. Box 97013
Redmond, WA 98073-9713
(425) 882-3700

ISBN 1-882588-52-5
ISBN 1-882588-54-1
P/N 061-0744-00